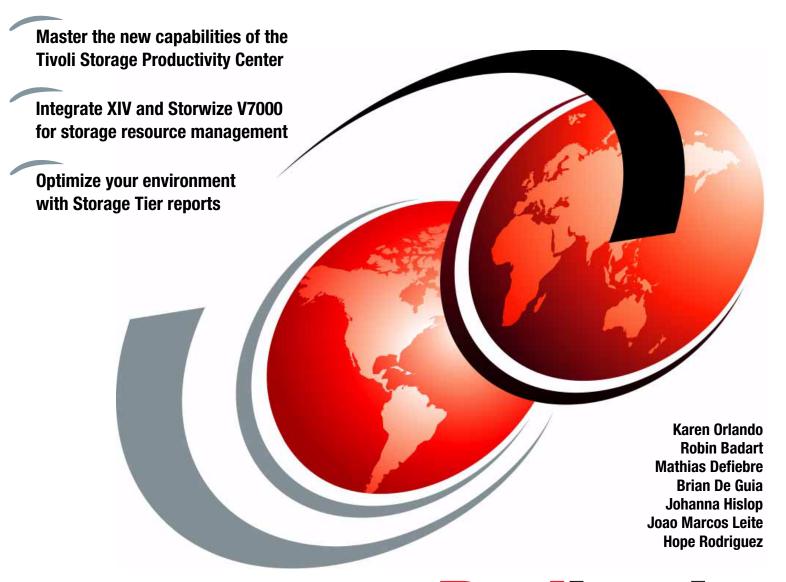


IBM Tivoli Storage Productivity Center V4.2 Release Guide



Redbooks



International Technical Support Organization

IBM Tivoli Storage Productivity Center V4.2 Release Guide

December 2011

Note: Before using this information and the product it supports, read the information in "Notices" on page xiii.
Second Edition (December 2011)
This edition applies to Version 4, Release 2 of IBM Tivoli Storage Productivity Center (product numbers 5608-WB1, 5608-WB2, 5608-WB3, 5608-WC3, 5608-WC4,5608-E14).

Contents

Notices	
Trademarks	XIV
Preface	xv
The team who wrote this book	
Now you can become a published author, too!	. xviii
Comments welcome	
Stay connected to IBM Redbooks	xix
Chapter 1. Tivoli Storage Productivity Center V4.2 introduction and overview	1
1.1 Introduction and functions of IBM Tivoli Storage Productivity Center	
1.2 Architecture	
1.2.1 Architecture overview	
1.2.2 Data server	4
1.2.3 Device server	5
1.2.4 Tivoli Integrated Portal	5
1.2.5 Tivoli Storage Productivity Center for Replication	5
1.2.6 DB2 database	
1.2.7 Agents	6
1.2.8 Interfaces	6
1.2.9 Integration with other applications	7
1.3 Tivoli Storage Productivity Center family	
1.3.1 Tivoli Storage Productivity Center for Data	
1.3.2 Tivoli Storage Productivity Center for Disk	
1.3.3 Tivoli Storage Productivity Center for Disk Select	
1.3.4 Tivoli Storage Productivity Center Select	
1.3.5 Tivoli Storage Productivity Center Basic Edition	
1.3.6 Tivoli Storage Productivity Center Standard Edition	
1.3.7 Tivoli Storage Productivity Center for Replication	
1.3.8 IBM System Storage Productivity Center Version 1.5	
1.3.9 IBM System Director Storage Control	
1.4 New functions since version Tivoli Storage Productivity Center 4.1	
1.4.1 New for Tivoli Storage Productivity Center Version 4.2.0	
1.4.2 New for Tivoli Storage Productivity Center Version 4.2.1	
1.4.3 New for Tivoli Storage Productivity Center Version 4.2.2	
1.5 New functions since version Tivoli Storage Productivity Center for Replication 4.1.1.5.1 New for IBM Tivoli Storage Productivity Center for Replication 4.2.0	
1.5.2 New for IBM Tivoli Storage Productivity Center for Replication 4.2.1	
1.5.3 New for IBM Tivoli Storage Productivity Center for Replication 4.2.2	
1.5.3 New for IBM Tivoli Storage Productivity Center for Replication 4.2.2	10
Chapter 2. Tivoli Storage Productivity Center installation on Windows	21
2.1 Tivoli Storage Productivity Center installation	
2.1.1 Installation overview	
2.1.2 Product code media layout and components	
2.2 Preinstallation steps for Windows	
2.2.1 Verifying system hardware and software prerequisites	
2.2.2 Verifying primary domain name systems	
2.2.3 Activating NetBIOS settings	
2.2.4 User IDs and passwords to be used and defined	28

2.3 Installing the prerequisite for Windows 2.3.1 DB2 installation 2.3.2 Verifying the installation 2.4 Installing Tivoli Storage Productivity Center components 2.4.1 Creating the Database Schema 2.4.2 Installing Tivoli Storage Productivity Center components 2.4.3 Agent installation. 2.4.4 Disabling Tivoli Storage Productivity Center or Tivoli Storage Productivity Center Replication	. 28 . 37 . 38 . 39 . 46 . 66 r for . 67
2.5 Applying a Tivoli Storage Productivity Center Upgrade	. 70
Chapter 3. Tivoli Storage Productivity Center installation on Linux	
3.1 Tivoli Storage Productivity Center installation on Linux	
3.1.2 Product code media layout and components	
3.2 Preinstallation steps for Linux	
3.2.1 Verifying system hardware and software prerequisites	
3.2.2 Prerequisite component for Tivoli Storage Productivity Center V4.2	. 82
3.2.3 DB2 installation: Using the GUI installation program	
3.2.4 Verifying that DB2 is installed correctly	
3.3 Installing Tivoli Storage Productivity Center components	
3.3.1 Preparing for the installation	
3.3.2 Creating the database schema	
3.3.3 Installing Tivoli Storage i Toddetivity Center Servers, Got, and General Servers	107
Chapter 4. Tivoli Storage Productivity Center installation on AIX	125
4.1 Tivoli Storage Productivity Center installation on AIX	
4.1.1 Installation overview	
4.1.2 Product code media layout and components	
4.2 Preinstallation steps for AIX	
4.2.1 Verifying system hardware prerequisites	
4.2.2 Verifying system software prerequisites	
4.3. Installing the prerequisite for AIX	
4.3.1 DB2 installation: Using the command line	
4.3.2 Verifying that DB2 is installed correctly	
4.4 Installing Tivoli Storage Productivity Center components	
4.4.1 Preparing for the installation	
4.4.2 Creating the database schema	135
4.4.3 Installing Tivoli Storage Productivity Center components	142
Chantas E. Migrating Tivali Staroga Draductivity Contas have and to august level	161
Chapter 5. Migrating Tivoli Storage Productivity Center base code to current level 5.1 Migration considerations	
5.1.1 Prerequisites	
5.1.2 Database considerations	
5.1.3 Tivoli Storage Productivity Center for Replication considerations	163
5.2 Credentials migration tool	170
5.2.1 Migration methods	
5.2.2 Running the credentials migration tool before the upgrade	
5.2.3 Running the credentials migration tool during the upgrade	
5.2.4 Running the credentials migration tool after the upgrade	
5.3 Agent Manager, Data agent, Fabric agent considerations	
5.4.1 Migration from Tivoli Storage Productivity Center V3.x	
I mg. a	

5.4.2 Migration from Tivoli Storage Productivity Center V4.1 5.5 Upgrading Storage Resource Agent	. 200 . 201 . 203 . 206
5.7 Upgrading System Storage Productivity Center (SSPC)	
Chapter 6. Agent migration and upgrade 6.1 Common Agent Services and Storage Resource agent history 6.2 Benefits of Data agent and Fabric agent migrated to Storage Resource agent 6.3 Prerequisites 6.4 Scenarios to migrate from Common Agent Services to Storage Resource agent 6.4.1 Installation wizard 6.4.2 Tivoli Storage Productivity Center user interface 6.4.3 Command-line interface 6.5 CIMOM to NAPI	. 210 . 211 . 211 . 212 . 213
Chapter 7. Device configuration: Introducing the Native API	
7.1.1 Full list of changes	
7.1.2 Changed panels and tasks	. 224
7.1.3 Behind the scenes: the External Process Manager	
7.2 Solution design for device access	
7.2.1 Planning for NAPI and NAPI discovery	
7.2.2 Planning for CIMOM discovery	
7.2.3 Planning for the Configure Devices Wizard	
7.2.4 Planning for Monitoring Groups	
7.2.5 Planning for events	
7.3 Using the Configure Devices wizard	
7.3.1 Scenario	
7.3.2 Adding or configuring an IBM SAN Volume Controller (SVC) or Starwing V700	
7.3.3 Adding or configuring an IBM SAN Volume Controller (SVC) or Storwize V700 7.3.4 Adding or configuring an IBM XIV	
7.3.5 Adding or configuring a CIMOM	
7.3.5 Adding or configuring fabrics and switches	
7.4.1 Adding or configuring a Brocade Fabric or switch	
7.4.2 Adding or configuring a McDATA Fabric or switch	
7.4.3 Adding or configuring a Cisco Fabric or switch	
7.4.4 Adding or configuring a QLogic Fabric or switch	
7.4.5 Adding or configuring a mixed Brocade/McDATA Fabric	
7.5 Other enhancements and changes	
7.5.1 Help panels	
7.5.2 Deleting a device	
7.5.3 Alert log	
7.5.4 Job history	
7.5.5 Job Management	
Chapter 8. Storage Resource Agent	. 263
8.1 Overview	
8.2 Storage Resource Agent requirements	
8.2.1 User requirements	
8.2.2 Platform dependencies	. 266

8.2.3 Communication requirements and types	
8.3 Storage Resource Agent installation methods	
8.3.1 Local graphical installer	
8.3.2 Storage Resource Agent deployment from GUI	
8.3.3 Local or CLI installation of the Storage Resource Agent	
8.4 Database monitoring with Storage Resource Agent	282
8.4.1 Registering the database	283
8.4.2 Setting up probes and scans	285
8.4.3 Database capacity reports	289
8.4.4 Database usage reports	
8.5 IBM Power VM Virtual I/O Server (VIOS) support	293
8.5.1 VIOS LUN correlation supported environments	293
8.5.2 Installing the Storage Resource Agent on VIOS	294
8.5.3 Uninstalling the Storage Resource Agent on VIOS	294
8.5.4 Overview of new functionality available with SRA on VIOS	295
8.6 NetApp and N series monitoring	297
8.6.1 Overview of NAS support	297
8.6.2 Configuring Windows Storage Resource Agent	300
8.6.3 Configuring the UNIX Storage Resource Agent	
8.6.4 Retrieving and displaying data about the NAS filer	322
8.7 VMware support	331
8.8 VMware virtual machine reporting	331
8.9 Batch reporting	
8.10 Storage Resource Agent fabric function	
8.10.1 HBA library requirements	340
8.10.2 Storage Resource Agent fabric enhancements	
8.10.3 Fabric agent assignment	
8.11 Agent resource utilization	
8.12 HBA information reports	
8.13 Collecting Storage Resource Agent support data	
8.14 Clustering support	
Chapter 9. Tivoli Storage Productivity Center Select and Tivoli Storage Productivity	
Center for Disk Select	
9.1 Overview	
9.2 Supported devices and firmware levels	
9.3 Licensing methodology	
9.4 Key benefits	347
Chapter 10. Tivoli Storage Productivity Center for Replication	240
10.1 New V4.2 features and functions	
10.2 Open HyperSwap replication	
10.2.1 Description	
10.2.2 Prerequisites	
10.2.3 Setting up Open HyperSwap session	
10.2.4 Performing Open HyperSwap	
• • • • • • • • • • • • • • • • • • • •	300
10.2.5 Tivoli Storage Productivity Center for Replication high availability with Open	272
HyperSwap	
· · · · · · · · · · · · · · · · · · ·	
10.3.1 Copy set soft removal scenario.	
10.3.2 Suspending the session after removing a copy set	
10.3.3 Removing a copy set from an Open HyperSwap	.,0.
10.4 Downloading log packages with the GUI	

10.6 SVC and Storwize V7000 enhancements	
10.6.1 SVC and Storwize V7000 space-efficient volumes	
10.6.2 SVC and Storwize V7000 incremental FlashCopy	
10.7.1 DS8000 extent space efficient volumes	
10.7.2 Global Mirror session enhancements	
10.7.3 Multiple Global Mirror sessions	
10.8 XIV enhancements	
10.8.1 Overview	
10.8.2 Snapshot sessions	
10.8.3 Metro Mirror Failover/Failback sessions	
10.8.4 Global Mirror Failover/Failback sessions	
10.8.5 Troubleshooting	
Chapter 11 VIV compart	407
Chapter 11. XIV support	
11.2 Adding XIV systems to Tivoli Storage Productivity Center	
11.3 XIV performance metrics	
11.4 XIV storage provisioning	
11.5 XIV and SAN Planner	
11.6 XIV and Storage Optimizer	
<u> </u>	
Chapter 12. SAN Planner	
12.1 Purpose of SAN Planner	
12.2 New features and functions	
12.3 Prerequisites for using SAN Planner	
12.4 Supported storage subsystems in SAN Planner	
12.4.1 Performance-based planning	
12.4.2 Space Only planning	
12.4.3 Replication planning	
12.5 Storage Resource Groups Monitoring and Alerting	
12.6 Creating a Space Only SAN Planner recommendation	
12.7.1 Introduction	
12.7.2 Disaster Recovery Planner scenario.	
12.7.2 Disaster Necovery Flamler Scenario	
12.9 SAN Planner with XIV replication	
12.9.1 Replication plan for XIV Snapshot	
12.9.2 Replication plan for XIV Metro and Global Mirror Failover/Failback	
Chapter 13. Job Management panel	
13.1 Background	
13.2 Job Management terms	
13.2.1 Schedule	
13.2.2 Run	
13.2.3 Job	
13.3 Job Management changes	
13.3.1 Default jobs	
13.3.2 CLI and event-driven jobs	
13.4 Job Management panel explained	
13.4.1 Entities details	
13.4.2 Schedules details	
13.4.3 Runs and Jobs details	
10.0 LAURIDIO	J+1

13.5.1 List Schedule log files	
Chapter 14. Fabric enhancements. 14.1 Fibre Channel over Ethernet (FCoE) support. 14.2 Additional switch models supported. 14.2.1 Brocade 3016, 5410, 5470, 5480, 7800, and M5440. 14.2.2 Brocade Silk Worm 7800 (IBM 2005-R06). 14.2.3 Brocade DCX-4S Backbone. 14.2.4 Brocade 8000. 14.2.5 Cisco Nexus 5000 Series. 14.3 Additional HBA and CNA models supported. 14.4 Integration with Brocade Data Center Fabric Manager and Brocade Network Advise 558	556 556 556 556 556 556 557
 14.4.1 Supported functions	559
Chapter 15. Reporting. 15.1 Reporting overview 15.1.1 Tivoli Storage Productivity Center GUI 15.1.2 Batch reports. 15.1.3 Tivoli Storage Productivity Center CLI (TPCTOOL) 15.1.4 Reporting database views. 15.2 Generating reports 15.3 Reporting example 15.4 Storage Capacity Pyramid 15.4.1 Overview 15.4.2 Tiers 15.5 Capacity licensing 15.5.1 Tivoli Storage Productivity Center for Disk Licensing 15.5.2 Tivoli Storage Productivity Center for Data Licensing. 15.5.3 Tivoli Storage Productivity Center Standard Edition licensing	568 568 569 569 570 578 578 578 588 589 590
15.6 VDisk mirroring reporting 15.7 Storwize V7000 and SVC internal disk reports 15.7.1 Data Manager report 15.7.2 Disk Manager report 15.7.3 Troubleshooting Internal disk reports 15.8 Storage Tier reports 15.8.1 Tivoli Common Reporting and Storage Tier reports 15.8.2 Terminology 15.8.3 TCR installation and setup 15.8.4 Running Storage Tier reports 15.8.5 Analyzing Storage Tier reports: Case Studies 15.8.6 Considerations 15.8.7 Migrating from the STAR tool 15.8.8 Supported environments 15.8.9 Non-supported environments	594 598 599 599 603 604 604 605 606 621 633 639 647 647
Chapter 16. Tivoli Storage Productivity Center database backup on Windows 16.1 Before you start	650 650

16.4 Offline backup to Tivoli Storage Manager setup steps	
16.4.1 Adding new variables to Windows	
16.4.2 Configuring Tivoli Storage Manager option file and password	
16.4.3 Rebooting the Tivoli Storage Productivity Center server	
16.4.4 Creating an offline backup to Tivoli Storage Manager script	657
16.5 Online backup to Tivoli Storage Manager setup steps	658
16.5.1 DB2 parameter changes for archive logging to Tivoli Storage Manager	
16.5.2 Creating an online backup script for Tivoli Storage Manager	660
16.6 Online backup to a file system setup steps	661
16.6.1 Setting up DB2 archive logging to a file system	661
16.6.2 Creating an online backup script to a file system	664
16.7 Performing offline database backups	664
16.7.1 Performing an offline backup to a file system	665
16.7.2 Performing an offline backup to Tivoli Storage Manager	665
16.8 Performing online database backup	666
16.8.1 Performing an online database backup to Tivoli Storage Manager	666
16.8.2 Performing an online backup to a file system	
16.9 Other backup considerations	
16.10 Managing database backup versions	
16.10.1 Managing backup versions for a file system	
16.10.2 Managing archive log files on a file system	
16.10.3 Managing backup versions that you store in Tivoli Storage Manager	
16.11 Verifying a backup file	
16.12 Restoring Tivoli Storage Productivity Center database	
16.12.1 Restoring from offline backups	
16.12.2 Restoring from online backups	
16.12.3 Potential agent issues after the restore process	
16.13 Backup scheduling and automation	
16.13.1 Frequency of full TPCDB backups	
16.13.2 TPCDB backup automation	
10.10.2 11 ODD backup automation	005
Chapter 17. Tivoli Storage Productivity Center database considerations	687
17.1 Database tuning	
17.1.1 Setting DB2 variables	
17.1.2 Tuning the database manager	
17.1.3 Changing DB2 active logs directory	
17.2 Repository database sizing	
17.2.1 Introduction	
17.2.2 Storage subsystem performance data sizing	
17.3 Repository calculation templates	
17.3.1 Worksheet: Sizing SVC and Storwize V7000 performance collection	
17.3.2 Worksheet: Sizing performance collection for XIV, DSxxxx, and non-IBM	704
subsystems	705
17.3.3 Worksheet: Sizing SAN switch performance collection.	
17.3.4 Worksheet: Sizing Tivoli Storage Productivity Center for Data repository	
17.4 Preparing for database backups	
17.4.1 Scripts provided	
17.4.2 Database backup	
17.4.3 Database backup method considerations	/11
Chapter 18. Tivoli Storage Productivity Center database backup on Linux	712
18.1 Before you start	
18.2 Common backup setup steps	
	, , ,

18.3 Offline backup to file system setup steps	715
18.4 Offline backup to Tivoli Storage Manager setup steps	716
18.4.1 Adding new variables to Linux	717
18.4.2 Configuring Tivoli Storage Manager option file and password	
18.4.3 Restarting DB2	
18.4.4 Creating an offline backup to Tivoli Storage Manager script	
18.5 Online backup to Tivoli Storage Manager setup steps	
18.5.1 DB2 parameter changes for archive logging to Tivoli Storage Manager	
18.5.2 Creating an online backup script for Tivoli Storage Manager	
18.6 Online backup to a file system setup steps	
18.6.1 Setting up DB2 archive logging to a file system	
18.6.2 Creating an online backup script to a file system	
18.7 Performing offline database backups	
18.7.1 Performing an offline backup to a file system	
18.7.2 Performing an offline backup to Tivoli Storage Manager	
18.8 Performing online database backup	
18.8.1 Performing an online database backup to Tivoli Storage Manager	
18.8.2 Performing an online backup to a file system	
18.9 Other backup considerations	
18.10 Managing database backup versions	
18.10.1 Managing backup versions for a file system	
18.10.2 Managing archive log files on a file system	
18.10.3 Managing backup versions that you store in Tivoli Storage Manager	
18.11 Verifying a backup file	
18.12 Restoring Tivoli Storage Productivity Center database	
18.12.1 Restoring from offline backups	
18.12.2 Restoring from online backups	
18.12.3 Potential agent issues after the restore process	
18.13 Backup scheduling and automation	
18.13.1 Frequency of full TPCDB backups	
18.13.2 TPCDB backup automation	/4/
Chapter 19. Tivoli Storage Productivity Center database backup on AIX	7/10
19.1 Before you start	
19.2 Common backup setup steps	
19.3 Offline backup to file system setup steps	
19.4 Offline backup to Tivoli Storage Manager setup steps	
19.4.1 Adding new variables to AIX	
19.4.2 Configuring Tivoli Storage Manager option file and password	
19.4.3 Restarting DB2	
19.4.4 Creating an offline backup to Tivoli Storage Manager script	
19.5 Online backup to Tivoli Storage Manager setup steps	
19.5.1 DB2 parameter changes for archive logging to Tivoli Storage Manager	
19.5.2 Creating an online backup script for Tivoli Storage Manager	
19.6 Online backup to a file system setup steps	
19.6.1 Setting up DB2 archive logging to a file system	
19.6.2 Creating online backup script to file system	
19.7 Performing offline database backups	
19.7.1 Performing an offline backup to a file system	
19.7.2 Performing an offline backup to Tivoli Storage Manager	
19.8 Performing online database backup	
19.8.1 Performing an online database backup to Tivoli Storage Manager	
, , , , , , , , , , , , , , , , , , ,	765

19.9 Other backup considerations	
19.10 Managing database backup versions	
19.10.1 Managing backup versions for a file system	
19.10.2 Managing archive log files on a file system	
19.10.3 Managing backup versions that you store in Tivoli Storage Manager	
19.11 Verifying a backup file	
19.12 Restoring Tivoli Storage Productivity Center database	
19.12.1 Restoring from offline backups	
19.12.2 Restoring from online backups	
19.12.3 Potential agent issues after the restore process	
19.13 Backup scheduling and automation	
19.13.1 Frequency of full TPCDB backups	
19.13.2 TPCDB backup automation	782
Chapter 20. Lessons learned	783
20.1 User-defined properties for Fabrics and Switches	
20.2 IBM Software Support Lifecycle	
20.3 IBM Support Assistant	
20.4 Certificate errors in Windows Internet Explorer	
20.4.1 Preventing address certificate mismatch errors	
20.4.2 Managing new certificates	
20.5 Tivoli Storage Productivity Center support matrix	
20.6 DB2 hints	
20.6.1 SQL5005C System Error	
20.6.2 User ID to stop and start DB2	
20.7 Extracting the installation images on AIX	
20.8 Tivoli Common Reporting hints	
20.8.1 Linux libraries	
20.8.2 Physical Address Extension on Windows	
20.8.3 Browser back button	
20.9 Hints and tips	797
Appendix A. DB2 table space considerations	
Selecting an SMS or DMS table space	
Advantages of an SMS table space	
Advantages of a DMS table space	801
Appendix B. Worksheets	803
User IDs and passwords	
Server information	
User IDs and passwords for key files and installation	
LDAP information	
Storage device information	
IBM System Storage Enterprise Storage Server, DS6000, DS8000	
IBM DS3000, DS4000, DS5000 family products	
IBM SAN Volume Controller (SVC)	
IBM XIV Storage System (XIV)	
IBM Storwize V7000 Midrange Disk System (Storwize)	
Other storage device information	
Appendix C. Configuring X11 forwarding	
Preparing the display export	
Preparation of the AIX server	
Preparation of the Windows workstation and Xming installation	817

aunching an Xming X Window session	
Appendix D. LDAP considerations Diverview Dirganization Tivoli Integrated Portal and Tivoli Storage Productivity Center Preparing to implement LDAP Deciding whether to install with OS or with LDAP backupConfig and restoreConfig Stopping and restarting services: Order matters Troubleshooting authentication problems: Configuration files Secure LDAP Useful LDAP tips 83	34 34 35 35 36 37 38
Appendix E. Tivoli Common Reporting: Remote Installation	1 2 3
Related publications 84 BM Redbooks publications 84 Other publications 84 Online resources 84 Help from IBM 85	19 19 19
medave OF	

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

НАСМР™

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX 5L™ **IBM®** System p®

AIX® Netcool® System Storage DS® Coanos® **NetView®** System Storage® DB2 Universal Database™ Power Systems™ System x® DB2® POWER6® System z®

Tivoli Enterprise Console® developerWorks® POWER7®

DS4000® ProtecTIER® **Tivoli®** DS6000™ WebSphere® pSeries® DS8000® **RDN® XIV®** z/OS® Enterprise Storage Server® Redbooks®

Redpaper™ FlashCopy® zSeries®

Redbooks (logo) @® HyperSwap® Storwize®

The following terms are trademarks of other companies:

Itanium, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Snapshot, Network Appliance, FilerView, Data ONTAP, NetApp, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Preface

IBM® Tivoli® Storage Productivity Center V4.2 is a feature-rich storage management software suite. The integrated suite provides detailed monitoring, reporting, and management within a single console.

This IBM Redbooks® publication is intended for storage administrators and clients who are installing and using the features and functions in IBM Tivoli Storage Productivity Center V4.2. The information in the book can be used to plan for, install, and customize the components of Tivoli Storage Productivity Center in your storage infrastructure. This release contains several important functional enhancements:

- It includes support for new licenses for Tivoli Storage Productivity Center Select and for Tivoli Storage Productivity Center for Disk Select to provide full Tivoli Storage Productivity Center functionality to Midrange.
- Support for IBM Storwize® V7000 (Storwize V7000). Storwize V7000 is a hardware and software solution that provides unmatched performance, availability, advanced functions, and highly scalable capacity.
 - Also new is the capability for customers to generate and see IBM System Storage® Storwize V7000 (Storwize V7000) and IBM System Storage SAN Volume Controller (SVC) Internal Disks report. For Storwize V7000 systems, you can now drill down from the internal disks node to view the internal storage on the system or generate a Storage Subsystem Disks report to view the internal storage on the system.
- Support is provided for new IBM Tivoli Storage Productivity Center Storage Tier reports that help clients by providing them with qualitative data to make important decisions about storage tiering. This is done by utilizing estimated capability and actual performance data for IBM System Storage SAN Volume Controller (SVC) and IBM Storwize V7000 Storage Tier reports.
- Storage Resource agent: Now supports file level and database level storage resource management (SRM) reporting for a broad set of platforms.
- ► IBM XIV® Storage System (XIV) support has been updated, adding discovery, provisioning, and performance management, as well as new replication support for three new XIV sessions: Snapshot, Metro Mirror failover/failback, and Global Mirror Failover/Failback.
- Storage area network (SAN) configuration planning supports the following features:
 - Best practice provisioning of replication relationships, as well as basic provisioning of non-IBM storage systems, are supported.
 - The SAN Planner now supports workload profiles and replication planning for XIV systems.
- Open IBM HyperSwap® for the IBM AIX® environment delivers application failover (no disruption of application I/O) across a synchronous mirror distance.

Step-by-step procedures are provided to help you with tasks such as migrating to Storage Resource agents, using Native APIs, using SAN configuration planning functions, and maintaining your IBM DB2® (DB2) database repository.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Karen Orlando is a Project Leader at the International Technical Support Organization, Tucson Arizona Center. Karen has over 25 years in the I/T industry with extensive experience in open systems, Product Test, and Information and Software Development of IBM hardware and software storage. She holds a degree in Business Information Systems from the University of Phoenix and is Project Management Professional (PMP) certified since 2005.

Robin Badart is a Tivoli Level 2 Support Software Engineer and is one of the lead L2 support engineers on Tivoli Storage Productivity Center. He has a BS in Computer Science from Northern Arizona University, and has worked for IBM for 31 years. Robin joined the IBM General Products Division in Tucson, Arizona in 1980 working in electronic design automation support, joined IBM Microelectronics in Burlington, Vermont working in the ASIC Design Center from 1997-2002, and returned to Tucson in Tivoli Level 2 Support from 2002 to the present date.

Mathias Defiebre is a Lab Services IT Specialist from Germany. He joined IBM in 1998 and has worked extensively with storage products. He graduated from the University of Cooperative Education Mannheim with a German Diploma in Information Technology Management and a Bachelor of Science. His areas of expertise include developing and providing services for the IBM TS76xx ProtecTIER® deduplication products, TSM integration and healthcheck solutions, and also performing Proof of Concept, workshops, and benchmarks with focus on data protection. He is an IBM Certified Specialist for IBM TotalStorage Networking and Virtualization Architectures. Mathias previously co-authored an earlier version of the Tivoli Storage Productivity Center Redbooks publication.

Brian De Guia is a Software Test Specialist in the Tivoli Storage Development lab in Tucson Arizona. He has 12 years of experience in the IT field. He has worked at IBM for 4 years. His areas of expertise include Tivoli Storage Productivity Center, IBM Storage Disk Systems, various OS platforms, and fabric/switch logical zoning topologies. Brian has written extensively on IBM Tivoli Storage Productivity Center, as well as various IBM internal and training documents.

Johanna Hislop is a Software Engineer in Rochester, Minnesota. She joined IBM in 1997 and is one of the lead L3 support engineers for Tivoli Storage Productivity Center. Her areas of expertise include Tivoli Storage Productivity Center, Fibre Channel switches, and working with customers and vendors for each. Johanna has written extensively on the Tivoli Storage Productivity support website and has several articles as well as a blog on the IBM developerWorks® website. She has a BS in Computer Science from the University of Minnesota Duluth.

Joao Marcos Leite is an IT Specialist from Brazil who joined IBM in 2000. He has worked for the IBM Systems and Technology Group for more than 10 years in the field of storage solutions design for clients, and is currently working as an IBM Systems Software architect for Latin America focused on storage software. Joao studied at the Universidade Federal do Parana in Curitiba, Brazil, and is a graduated Technologist in Data Processing. His areas of expertise include storage virtualization and storage management, and he co-authored a previous update to the SAN Volume Controller Implementation Redbooks publication. He is also a member of the Technology Leadership Council Brazil (TLC-BR), and an affiliate of the IBM Academy of Technology.

Hope Rodriguez joined IBM in 2008 after studying at New Mexico State University; she holds her degree in Information Systems. Hope is a Software Test Engineer for Tivoli Storage Productivity Center / System Storage Productivity Center in Tucson, AZ. She is also an IBM Certified Deployment Professional on Tivoli Storage Productivity Center and has experience with a wide array of IBM storage products.



Figure 1 The Team: Joao, Robin, Karen, Johanna, Hope, Mathias, and Brian

Thanks to the following people for their contributions to this project:

Mary Lovelace Larry Coyne Sangam Racherla Ann Lund International Technical Support Organization

Thanks to The Global Overview(GO) Training Team:

Jay Calder
Mike Lamb
Raymond Swank
Roman Druker
Scott Trevett
Todd Gerlacht
Yan (Charlie) H. Chu
William Maloney
IBM Tivoli IBM Software Group, Development

Thanks to the Tivoli Storage Productivity Center Tier Report white paper authors:

Mike Lamb
Patrick Leahy
Balwant Rai
Jackson Shea
IBM Tivoli IBM Software Group, Development

Katherine Keanev

TIvoli Storage Productivity Center, Software Development, Project Manager

Sondra Ashmore

Tlvoli Storage Productivity Center, Software Development, Executive Project Manager

Xin Wang

TIvoli Storage Productivity Center, Software Development, Tivoli Product Manager, TPC/SSPC

John Wyatt Thuan Ngyuen IBM Tivoli Storage Productivity Center Development and Support

Thanks also to the authors of the previous editions of this book.

The following individuals were authors of the first edition, *Tivoli Storage Productivity Center V4.2 Release Guide*, published in April 2011:

Mary Lovelace
Alejandro Berardinelli
H. Antonio Vazquez Brust
Harsha Gunatilakai
Hector Hugo Ibarra
Danijel Paulin
Markus Standau

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

► Send your comments in an email to:

redbooks@us.ibm.com

► Mail your comments to:

IBM Corporation, International Technical Support Organization Dept. HYTD Mail Station P099 2455 South Road Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

► Find us on Facebook:

http://www.facebook.com/IBMRedbooks

► Follow us on Twitter:

http://twitter.com/ibmredbooks

► Look for us on LinkedIn:

http://www.linkedin.com/groups?home=&gid=2130806

► Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

► Stay current on recent Redbooks publications with RSS Feeds:

http://www.redbooks.ibm.com/rss.html

1

Tivoli Storage Productivity Center V4.2 introduction and overview

In this chapter, we introduce IBM Tivoli Storage Productivity Center, providing a high level technical overview of the product, and its architecture and base components. We discuss the following topics:

- ► Introduction to IBM Tivoli Storage Productivity Center
- Functions
- Architecture
- Product family
- ► New functions since Tivoli Storage Productivity Center V4.1
- ► Contents of this book

Product name: IBM TotalStorage Productivity Center (TPC) was renamed Tivoli Storage Productivity Center in Version 4.1.

1.1 Introduction and functions of IBM Tivoli Storage Productivity Center

The IBM Tivoli Storage Productivity Center suite of storage infrastructure management tools can help customers improve time to value. It can reduce the complexity of managing their storage environments by centralizing, simplifying, and optimizing storage tasks that are associated with storage systems, storage networks, replication services, and capacity management.

Tivoli Storage Productivity Center includes the following functions:

- Storage resource management (SRM):
 - Reporting of volumes and file systems on a server level
 - Reporting on network-attached storage (NAS) and NetWare file systems
 - Reporting of databases capacity and usage
 - Constraint and quota reporting
- Storage subsystem management:
 - Volume allocation and assignment (provisioning)
 - Asset reporting
 - Performance reporting
 - Storage Subsystem element management
- Fabric management:
 - Zoning
 - Asset reporting
 - Performance reporting
- Replication management
- Alerting

In addition to these basic functions, Tivoli Storage Productivity Center includes more advanced functions that provide a set of analytics functions such as the following items:

- ▶ Topology Viewer
- ► Data path explorer
- Configuration history
- ► Storage optimizer
- SAN planner
- ► Configuration analytics

The type of license you have determines the functions that are available to you.

Table 1-1 summarizes the functions provided in the Tivoli Storage Productivity Center.

Available functions in each Tivoli Storage Productivity Center license are shown here.

Table 1-1 Functions of Tivoli Storage Productivity Center licenses

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Administrative Services						
Device Discovery	Yes ^a	Yes ^a	Yes ^a	Yes	Yes	Yes
Data Sources	Yes ^b	Yes ^b	Yes ^b	^c Yes	Yes	Yes
Configuration	Yes ^d	Yes ^d	Yes ^d	Yes	Yes	Yes
IBM Tivoli Storage Productivity Center						
Configuration Utility	Yes	Yes	Yes	Yes	Yes	Yes
Configure Devices	Yes	Yes	Yes	Yes	Yes	Yes
Job Management	Yes	Yes	Yes	Yes	Yes	Yes
Reporting: Batch reports User ID reports System Reports Rollup Reports Data Source Reports	Yes ^e	Yes ^e	Yes ^e	Yes ^e	Yes ^e	Yes ^e
Topology Viewer	Yes	Yes	Yes	Yes	Yes	Yes

- a. You cannot run the following discovery jobs in the Basic, Disk, and Disk Select editions: Netware Filer, Windows Domain, NAS, SAN FS, and VMware VI Data Source.
- b. You cannot associate VMware VI Data Sources or TPC Servers with Tivoli Storage Productivity Center in the Basic, Disk, and Disk Select editions.
- c. Data Sources for Data Edition does not support TPC Server Probes.
- d. The following nodes under *Configuration* are not available in the Basic, Disk, and Disk Select editions:
 - License Keys
 - Quota and Constraint e-mail Address Rules
 - Scan/Probe Agent Administration
 - Manual NAS/Netware Server Entry
 - NetWare Tree Logins
 - · Resource History Retention for Databases
 - Removed Resource Retention for Databases
 - Configuration History Settings
- e. For more licensing details, including report availability for each license type, see Chapter 2, "Product Overview" in the *IBM Tivoli Storage Productivity Center Version 4.2.2 User's Guide*, SC27-2338.

1.2 Architecture

The IBM Tivoli Storage Productivity Center consists of several key components. In this section, we describe these components and how they are related. We also describe the interfaces that you can use to access Tivoli Storage Productivity Center, and its integration with other products.

1.2.1 Architecture overview

Figure 1-1 is an architectural overview of IBM Tivoli Storage Productivity Center.

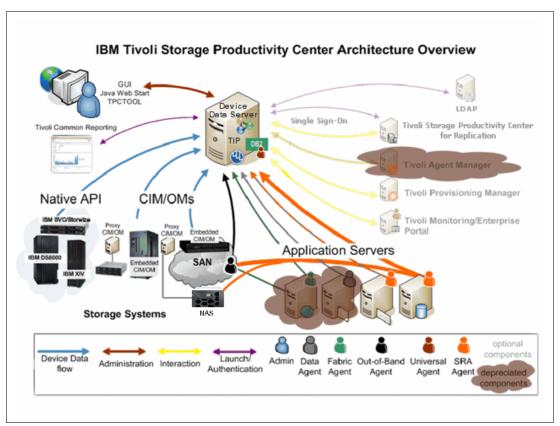


Figure 1-1 Overview of IBM Tivoli Productivity Center Architecture

1.2.2 Data server

This component is the control point for product scheduling functions, configuration, event information, reporting, and graphical user interface (GUI) support. It coordinates communication with, and data collection from, agents that scan file systems and databases to gather storage demographics and populate the database with results.

Automated actions can be defined to perform file system extension, data deletion, and Tivoli Storage Manager backup or archiving, or event reporting when defined thresholds are encountered. The Data server is the primary contact point for GUI user interface functions. It also includes functions that schedule data collection and discovery for the Device server.

1.2.3 Device server

This component discovers, gathers information, analyzes performance, and controls storage subsystems and SAN fabrics. It coordinates communication with and data collection from agents that scan SAN fabrics and storage devices.

1.2.4 Tivoli Integrated Portal

IBM Tivoli Storage Productivity Center V4 is integrated with IBM Tivoli Integrated Portal. This integration provides functionalities such as single sign-on and the use of Tivoli Common Reporting.

Single sign-on

With single sign-on, you can access Tivoli Storage Productivity Center and then Tivoli Storage Productivity Center for Replication, IBM System Storage DS8000®, SAN Volume Controller 6.1, and Storwize V7000 by using a single user ID and password.

Tivoli Common Reporting

Tivoli Common Reporting is a component provided by Tivoli Integrated Portal. It is one possible option to implement customized reporting solutions using SQL database access, providing output in HTML, PDF, or Microsoft Excel. For more detail, see Chapter 15, "Reporting" on page 567.

Tivoli Common Reporting is intended to provide a platform to reproduce custom reports in an easy way or for reports that are to be run repeatedly (typically daily, weekly, or monthly).

Reports: The new Tivoli Storage Productivity Center Tier reports require Tivoli Common Reporting to be configured. Use Storage Tier reports to help determine the best Storage Tier configuration for the needs of your environment. You can use the data that IBM Tivoli Storage Productivity Center collects from IBM System Storage SAN Volume Controller and IBM Storwize V7000 managed-disk groups (MDiskgroups) to create and run Storage Tier performance reports on those managed-disk groups and related VDisks.

For details on Storage Tier reports, see Chapter 15, "Reporting" on page 567.

1.2.5 Tivoli Storage Productivity Center for Replication

Starting with Tivoli Storage Productivity Center V4.1, the IBM Tivoli Storage Productivity Center for Replication product is integrated into Tivoli Storage Productivity Center. Currently the integration is limited to basic functions such as providing launch-in-context links in the Tivoli Storage Productivity Center GUI. It crosschecks when a volume is deleted with Tivoli Storage Productivity Center and mapping of user roles.

1.2.6 DB2 database

A single database instance serves as the repository for all Tivoli Storage Productivity Center components. This repository is where all of your storage information and usage statistics are stored. All agent and user interface access to the central repository is done through a series of calls and requests made to the server. All database access is done using the server component to maximize performance and to eliminate the need to install database connectivity software on your agent and UI machines.

1.2.7 Agents

Outside of the server, several interfaces can be used to gather information about the environment. The most important sources of information are the Tivoli Storage Productivity Center agents (Storage resource agent, Data agent and Fabric agent) and Storage Management Initiative Specification (SMI-S) enabled storage devices that use a CIMOM agent (either embedded or as a proxy agent).

Storage Resource agent, CIM agents, and out-of-band fabric agents gather host, application, storage system, and SAN fabric information and send that information to the Data server or Device server.

Attention: Data agents and Fabric agents are supported in Tivoli Storage Productivity Center V4.2. However, no new functions were added to those agents for this release. For optimal results when using Tivoli Storage Productivity Center, migrate the Data agents and Fabric agents to Storage Resource agents.

1.2.8 Interfaces

As Tivoli Storage Productivity Center gathers information from your storage (servers, subsystems, and switches) across your enterprise, it accumulates a repository of knowledge about your storage assets and how they are used. You can use the reports provided in the user interface view and analyze that repository of information from various perspectives to gain insight into the use of storage across your enterprise.

With the user interfaces (UI), users can request information, and then generate and display reports based on that information. Certain user interfaces can also be used for configuration of Tivoli Storage Productivity Center or storage provisioning for supported devices.

The following interfaces are available for Tivoli Storage Productivity Center:

► Tivoli Storage Productivity Center GUI:

This interface is the central point of Tivoli Storage Productivity Center administration. Here you have the choice of configuring Tivoli Storage Productivity Center after installation, define jobs to gather information, initiate provisioning functions, view reports, and work with the advanced analytics functions.

► Java Web Start GUI:

When you use Java Web Start, the regular Tivoli Storage Productivity Center GUI is downloaded to your workstation and started automatically, so you do not have to install the GUI separately. The main reason for using the Java Web Start is that it can be integrated into other products (for example, Tivoli Integrated Portal). By using launch-in-context from those products, you are guided directly to the select panel. The launch-in-context URLs can also be assembled manually and be used as bookmarks.

► TPCTOOL:

This command-line interface (CLI) program interacts with the Tivoli Storage Productivity Center Device server. Most frequently it is used to extract performance data from the Tivoli Storage Productivity Center repository database to create graphs and charts with multiple metrics, with various unit types and for multiple entities (for example, Subsystems, Volumes, Controller, Arrays) using charting software. Commands are entered as lines of text (that is, sequences of types of characters) and output can be received as text. Furthermore, the tool provides queries, management, and reporting capabilities, but you cannot initiate discoveries, probes and performance collection from the tool.

► Database access:

Starting with Tivoli Storage Productivity Center V4, the Tivoli Storage Productivity Center database has views to access to the data stored in the repository, which allows you to create customized reports. The views and the required functions are grouped together into a database schema called TPCREPORT. To create customized reports using the TPCREPORT database schema, you need to have sufficient knowledge about SQL. To access the views, DB2 supports various interfaces, for example, JDBC and ODBC.

1.2.9 Integration with other applications

In this section, we describe Tivoli Storage Productivity Center integration with IBM Tivoli Storage Manager for backup or archival of files, IBM Tivoli Enterprise Console®, OMNIbus, or any other Simple Network Management Protocol (SNMP) manager for alert notification.

Integration with IBM Tivoli Storage Manager

Use the Archive/Backup function available within the Reporting facility to define IBM Tivoli Storage Manager archive or backup jobs to run against the files that you select from reports. With this function, you can select a specific file or group of files from Data Manager reports that you want to archive or back up by using Tivoli Storage Manager. Several reports include largest files, most obsolete files, duplicate files, and constraint violations.

The results of the IBM Tivoli Storage Manager backup-archive commands are viewable through the graphical user interface (GUI). In the case of constraints configured to archive-backup violating files, the results are included in the agent scan job logs (scans are responsible for enforcing constraints). In the case of file-report-driven archive-backup operations, a new type of job (archive-backup job) is created. The results of the backup operations in this case are found in archive-backup job logs.

Simple Network Management Protocol (SNMP)

For users planning to use the SNMP trap alert notification capabilities of Tivoli Storage Productivity Center, SNMP Management Information Base (SNMP MIB) files are included on the installation media. The MIB is provided for use by your SNMP management console software (for example, IBM Tivoli NetView® or HP OpenView). As a result, you can better view Tivoli Storage Productivity Center-generated SNMP traps from within your management console software.

Integration with Tivoli Enterprise Console Netcool/OMNIbus

Tivoli Storage Productivity Center can use the Event Integration Facility (EIF) to send messages to the IBM Tivoli Enterprise Console or the follow-on product, Netcool/OMNIbus. This way can allow one of the two central monitoring applications to consider Tivoli Storage Productivity Center alerts in causal analysis for problems. Tivoli Enterprise Console IBM Netcool® OMNIbus is added as a destination for alerts, in addition to SNMP Trap and Windows Event Log.

Integration with Tivoli Business Service Manager

Tivoli Storage Productivity Center can help integrate Storage resources (that is, Fabrics, Switches, Storage subsystems, and Servers) into Tivoli Business Service Manager (TBSM) Services. Tivoli Storage Productivity Center resources can be imported into TBSM using Tivoli Storage Productivity Center Discovery Library Adapter (DLA) output or via Tivoli Application Dependency Discovery Manager (TADDM) TPCStorage sensor.

Further, availability and status alerts generated in TPC can be routed to TBSM via Onmnibus EIF probe. TBSM views can correlate these events to TPC Storage resources and display corresponding status change to notify Business service clients of changes in the environment.

1.3 Tivoli Storage Productivity Center family

In this section, we describe the Tivoli Storage Productivity Center family of products.

1.3.1 Tivoli Storage Productivity Center for Data

Tivoli Storage Productivity Center for Data provides over 400 enterprise-wide reports, monitoring and alerts, policy-based action and file-system capacity automation in a heterogeneous environment. Tivoli Storage Productivity Center for Data can help improve capacity usage of file systems and databases and add intelligence to data protection and retention practices.

1.3.2 Tivoli Storage Productivity Center for Disk

Tivoli Storage Productivity Center for Disk is designed to provide storage device configuration and management from a single console. It includes performance capabilities to help monitor and manage performance, and measure service levels by storing received performance statistics into database tables for later use. Policy-based automation enables event action based on business policies. It sets performance thresholds for the devices based on selected performance metrics, generating alerts when those thresholds are exceeded. Tivoli Storage Productivity Center for Disk helps simplify the complexity of managing multiple SAN-attached storage devices.

1.3.3 Tivoli Storage Productivity Center for Disk Select

Tivoli Storage Productivity Center for Disk Select, formerly known as Tivoli Storage Productivity Center for Disk Midrange Edition, can help reduce the complexity of managing storage devices by allowing administrators to configure, manage and monitor performance of their entire storage infrastructure from a single console.

Tivoli Storage Productivity Center for Disk Select provides the same features and functions as does Tivoli Storage Productivity Center for Disk. However, it is limited to managing IBM XIV Storage System (XIV), IBM System Storage DS3000 (DS3000), IBM System Storage DS4000® (DS4000), and IBM System Storage DS5000 (DS5000) devices. It provides performance management, monitoring, and reporting for these devices.

1.3.4 Tivoli Storage Productivity Center Select

This Select Edition includes the same functions and features available in IBM Tivoli Storage Productivity Center Standard Edition. It includes advanced analytics for storage provisioning, change management, and performance optimization capabilities. It also offers additional management, control, and performance reporting for the Fibre Channel SAN infrastructure.

These features and functions are available for IBM System Storage DS3000 (DS3000), IBM System Storage DS4000 (DS4000), IBM System Storage DS5000 (DS5000), IBM System Storage SAN Volume Controller (SVC), IBM System Storage Storwize V7000 (Storwize V7000), and IBM XIV Storage System (XIV).

1.3.5 Tivoli Storage Productivity Center Basic Edition

IBM Tivoli Storage Productivity Center Basic Edition provides basic device management services for IBM System Storage DS3000 (DS3000), IBM System Storage DS4000 (DS4000), IBM System Storage DS5000 (DS5000), IBM Stem Storage DS6000[™] (DS6000), IBM Stem Storage DS8000 (DS8000), IBM XIV Storage System (XIV), IBM System Storage SAN Volume Controller (SVC), IBM System Storage Storwize V7000 (Storwize V7000), and heterogeneous storage environments.

Storage administrators can use this tool as a simple way to conduct device management for multiple storage arrays and SAN fabric components from a single integrated console. IBM Tivoli Storage Productivity Center Basic Edition also does discovery and asset management of tape libraries, specifically IBM 3494 and 3584 Tape Libraries.

1.3.6 Tivoli Storage Productivity Center Standard Edition

IBM Tivoli Storage Productivity Center Standard Edition is one of the industry's most comprehensive storage resource management solutions by combining the consolidated benefits of the four previous components as one bundle.

In addition to the benefits and features of Data, Disk, Disk Select and Basic Edition, IBM Tivoli Storage Productivity Center Standard Edition includes advanced analytics for storage provisioning, change management and performance optimization capabilities. It also offers additional management, control and performance reporting for the Fibre Channel SAN infrastructure.

1.3.7 Tivoli Storage Productivity Center for Replication

IBM Tivoli Storage Productivity Center for Replication helps to manage the advanced copy services provided by the IBM Enterprise Storage Server® Model 800, IBM System Storage DS8000 (DS8000), IBM System Storage DS6000 (DS6000), IBM System Storage SAN Volume Controller (SVC), IBM System Storage Storwize V7000 (Storwize V7000), and IBM XIV Storage Systems (XIV).

Licensing: Tivoli Storage Productivity Center for Replication comes with a basic license for One Site Business Continuity licensing. When you install Tivoli Storage Productivity Center for Replication, there are several licenses that you can purchase. For details on Two Site and Three Site Business Continuity licensing, see Chapter 1, "Planning for the IBM Tivoli Storage Productivity Center family" in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Version 4.2.2, Installation and Configuration Guide*, SC27-2337-05.

1.3.8 IBM System Storage Productivity Center Version 1.5

The IBM System Storage Productivity Center (SSPC) is a hardware appliance that consolidates IBM storage administration and configuration utilities into a single console.

The new features, functions and enhancements that are included in IBM System Storage Productivity Center Version 1.5 (SSPC 1.5) are as follows:

- ► Machine type MC5
- Tivoli Storage Productivity Center 4.2.1:

IBM Tivoli Storage Productivity Center Basic Edition 4.2.1 and IBM Tivoli Storage Productivity Center for Replication Basic Edition 4.2.1 are pre installed on the System Storage Productivity Center server.

- ▶ Microsoft Windows Server 2008 R2 Standard operating system for 64-bit processors
- Optional pre installed host bus adapter card:

To provide additional storage on the SSPC server, SSPC V1.5 offers an optional host bus adapter (HBA). Thus you can move the Tivoli Storage Productivity Center database from the SSPC server to an internal redundant hard disk drive or to an IBM System Storage DS8000.

- Documentation about international power requirements
- ▶ DB2 9.7
- ▶ IBM System Storage DS8000 (DS8000)Release 6.0
- ► SAN Volume Controller Console (SVC) Release 6.1:

SSPC supports IBM System Storage SAN Volume Controller 6.1, but the software is no longer pre installed on the SSPC server. Instead, you can start the console from the web browser on the SSPC desktop.

► IBM Storwize V7000 (Storwize V7000):

Storwize V7000 is a hardware and software solution that provides unmatched performance, availability, advanced functions, and highly scalable capacity. Storwize V7000 offers IBM storage virtualization, SSD optimization and "thin provisioning" technologies built in to improve storage utilization.

► IBM System Storage DS® Storage Manager Release 10.70:

The IBM System Storage DS Storage Manager user interface is available for you to optionally install on the SSPC server or on a remote server.

► IBM Java Release 1.6:

IBM Java 1.6 is pre installed and can be used with DS Storage Manager 10.70. You do not need to download Java from Oracle.

▶ DS CIM agent management commands:

The DS CIM agent management commands (DSCIMCLI) for Release 6.0 are pre installed on the SSPC server.

Optional media to recover image for 2805-MC5

For further information about SSPC 1.5, see the IBM Redbooks publication; *IBM System Storage Productivity Center Deployment Guide*; SG24-7560.

1.3.9 IBM System Director Storage Control

The IBM System Director Storage Control is designed to be used only as an embedded version of Tivoli Storage Productivity Center, without a GUI, only under a consuming application. At the time of the writing of this book, IBM System Director Storage Control is only used under IBM System Director Version 6.2.1 using Tivoli Storage Productivity Center V4.2.1 and higher.

1.4 New functions since version Tivoli Storage Productivity Center 4.1

This section lists new functions in Tivoli Storage Productivity Center since V4.1 and is divided into three segments;

- ▶ New for Tivoli Storage Productivity Center Version 4.2.0
- ▶ New for Tivoli Storage Productivity Center Version 4.2.1
- ▶ New for Tivoli Storage Productivity Center Version 4.2.2

1.4.1 New for Tivoli Storage Productivity Center Version 4.2.0

This section summarizes new functions in Tivoli Storage Productivity Center since V4.1.

Native Storage system interfaces provided for DS8000, SVC and XIV

To improve the management capabilities and performance of data collection for the DS8000, SAN Volume Controller (SVC), and XIV storage systems, native storage system interfaces are provided. Now Tivoli Storage Productivity Center communicates with these storage systems through ESSNI interface for DS8000, SSH for SVC and XCLI for XIV. These interfaces replace the CIM agent (SMI-S agent) implementation.

SAN Volume Controller

When you add the SAN Volume Controller to Tivoli Storage Productivity Center, you must supply a private SSH key. Tivoli Storage Productivity Center requires an OpenSSH key or PuTTY (.ppk) key format. More information is given in "IBM SAN Volume Controller (SVC) or Storwize V7000" on page 230.

Configure Devices wizard

Use the Configure Devices wizard to set up storage devices for monitoring by IBM Tivoli Storage Productivity Center. The wizard guides you through the steps for adding a device as a data source, running a discovery, including devices in groups, specifying alerts, and setting up data collection schedules. The wizard supports configuration of storage subsystems, fabrics and switches, computers and tape libraries.

Job Management panel

IUse the Job Management panel in the user interface to view and manage the schedules, runs, and jobs that are related to the storage entities that are monitored by Tivoli Storage Productivity Center.

Storage Resource agents

The Storage Resource agents now perform the functions of the Data agents and Fabric agents. (Out-of-band Fabric agents are still supported and their function has not changed.)

Before you migrate an existing Data agent or Fabric agent to a Storage Resource agent or deploy a new Storage Resource agent, make sure that the product functions you want to use on the monitored devices are available for those agents.

Data agents, Fabric agents, and Agent Manager

The Data agents, Fabric agents, and Agent Manager are supported in the current release but no new functions were added to these components. The legacy Data agents and Fabric agents are supported at Version 3.3.x and 4.1.x, and can communicate with the Tivoli Storage Productivity Center 4.2 server. You can also select to migrate the Data agents and Fabric agents to Storage Resource agents.

The Tivoli Storage Productivity Center V4.2 installation program does not support installation of the Data agent or Fabric agent. If you want to install the legacy Data agent or Fabric agent, you must have a previous Tivoli Storage Productivity Center installation program that supports installing the Data agent or Fabric agent.

If you are installing DB2 9.7 and want to use the Agent Manager, you must install a new release of Agent Manager 1.4.2 or later. Agent Manager 1.3 does not support DB2 9.7.

Most of the information about Tivoli Common Agent Services has been removed from the Tivoli Storage Productivity Center V4.2 documentation. However, this information is still available in the Information Center for Tivoli Storage Productivity Center Version 4.1.1.

SAN Planner

Tivoli Storage Productivity Center provides a new SAN Planner wizard, which has been enhanced to support the following functions.

SAN Volume Controller and Storwize V7000 with provisioning and workload profiles

The SAN Planner guidelines are limited to SAN Volume Controller and Storwize V7000 front-end operations only. The support includes the creation and provisioning of VDisks with the recommended I/O group and preferred node for each VDisk. The SAN Planner does not support back-end operations such as the creation of new MDisks or the creation or expansion of MDisk groups.

Space Only planning for all storage subsystems

All storage subsystems supported by Tivoli Storage Productivity Center can be used for space-based planning.

Resiliency profile for Tivoli Storage Productivity Center for Replication

The SAN Planner has a new profile called the resiliency profile for resilient resources. The resiliency profile is created internally when you select different options in the SAN Planner wizard. Resilient resource planning is available only for devices which are supported by Tivoli Storage Productivity Center for Replication. The supported devices are DS8000, DS6000, Enterprise Storage Server, SAN Volume Controller, and Storwize V7000.

Space-efficient volumes

The SAN Planner now has an option to provision space-efficient volumes on supported storage subsystems. These storage subsystems are: SVC (v4.3 or later), Storwize V7000, XIV (v10.1 or later), and DS8000 (v4.3 or later).

Encrypted volumes

Tivoli Storage Productivity Center supports encrypted volumes for the DS8000. The SAN Planner has been enhanced to allow input from the user for encrypted volumes as needed. The SAN Planner currently supports encrypted volumes for the DS8000 and SAN Volume Controller and Storwize V7000 (if the DS8000 is used as a back-end device).

Candidate storage resource group

For the SAN Planner, the candidate storage resource group is a container of storage resources. When you provide a candidate storage resource group for input, the SAN Planner uses the storage subsystems, pools, and volumes from that storage resource group for provisioning new storage.

New switches supported in toleration mode only

Tivoli Storage Productivity Center displays both Fibre Channel over Ethernet and FC ports in the switch port lists. Now supports the following switches: Brocade 8000, Brocade DCX-4S Backbone and Cisco Nexus 5000.

Note that not all functions are supported, for example it does not support the Converged Enhanced Ethernet (CEE) or Fibre Channel over Ethernet (FCoE) connectivity functions.

Brocade Data Center Fabric Manager

Tivoli Storage Productivity Center supports the new embedded SMI Agent in the Data Center Fabric Manager (DCFM) 10.4.0 or later (it still supports separate non-embedded SMI Agent). This DCFM manages both the McDATA and Brocade switches. The DCFM manages multiple fabrics within and across data centers. When you configure DCFM, you set up one switch to be the "master switch" that interconnects to all the other switches in the fabric. The embedded SMI Agent supports the SMI-S 1.2 standards.

IBM Tivoli Storage Productivity Center Monitoring Agent

The IBM Tivoli Storage Productivity Center Monitoring Agent can be used by the IBM Tivoli Enterprise Monitoring Server to monitor systems in your enterprise. This agent is an optional program you can install and use in your enterprise.

IBM XIV Storage System

Tivoli Storage Productivity Center supports performance monitoring and provisioning for XIV storage systems through the native interface.

1.4.2 New for Tivoli Storage Productivity Center Version 4.2.1

This section describes the new products available.

IBM SAN Volume Controller(SVC) Version 6.1

Tivoli Storage Productivity Center supports SAN Volume Controller Version 6.1. Tivoli Storage Productivity Center also supports launch in context and single sign-on for SAN Volume Controller.

Support: Launch in context and single sign-on were not supported for SAN Volume Controller Version 5.1 or earlier.

IBM Storwize V7000 (Storwize V7000)

Storwize V7000 is a hardware and software solution that provides online storage optimization through real-time data compression. This solution helps to reduce costs without performance degradation.

New performance metrics, counters, and thresholds for DS8000 and SAN Volume Controller

The Performance Manager now collects, reports, and alerts users on various performance metrics related to fiber channel link problems in the SAN environment.

Some of the thresholds you can define are as follows:

- ► Error (illegal) frame rate for DS8000 ports
- ► Link failure rate for SAN Volume Controller and DS8000 ports
- ▶ Invalid CRC rate for SAN Volume Controller, DS8000, and switch ports
- Invalid transmission word rate for SAN Volume Controller, DS8000, and switch ports
- ► Zero buffer-buffer credit timer for SAN Volume Controller ports

XIV system enhancements

In Tivoli Storage Productivity Center 4.2, users can add an XIV storage system through a single IP address. However, XIV storage systems have the capability of defining up to three administrative nodes, each with their own IP address. In Tivoli Storage Productivity Center 4.2.1, if an XIV storage system is configured with multiple administrative nodes, Tivoli Storage Productivity Center detects the IP addresses for these nodes. If Tivoli Storage Productivity Center fails to connect to one of the IP addresses, then an attempt is made to connect to the XIV storage system using one of the other IP addresses.

LUN correlation support for Virtual SCSI disks on the Virtual I/O Client

The Virtual SCSI environment consists of a Virtual SCSI Server Adapter or Adapters that are created on the Virtual I/O Server and mapped to a Virtual I/O Client. If you have a Virtual SCSI environment and the virtual target device has a one-to-one mapping to a storage volume, Tivoli Storage Productivity Center can extract the storage system serial number and correlate the Virtual SCSI disk on the LPAR to a corresponding storage volume. This function means that the Storage Resource agent installed on the Virtual I/O Client (LPAR), can collect this information.

1.4.3 New for Tivoli Storage Productivity Center Version 4.2.2

This section describes the new products available.

Tivoli Storage Productivity Center Select

This Select Edition includes the same functions and features available in IBM Tivoli Storage Productivity Center Standard Edition:

- ► Disk, fabric, tape, and data management functions
- Storage system performance monitoring, fabric performance monitoring, chargeback, and database monitoring
- ► Analytical functions: Configuration Analysis, SAN Planner, Storage Optimizer, and Configuration History

This offering is designed to support IBM entry-level and midrange System Storage DS3000, System Storage DS4000, System Storage DS5000, SAN Volume Controller, Storwize V7000, and XIV.

Tivoli Storage Productivity Center for Disk Select

Tivoli Storage Productivity Center for Disk Select provides the same function as the previously offered Tivoli Storage Productivity Center for Disk Midrange Edition provided:

- ▶ Basic disk, fabric, tape, and data management functions
- ► Storage system performance monitoring, fabric performance monitoring, chargeback, and database monitoring
- ► Analytical functions: Configuration Analysis, SAN Planner, Storage Optimizer, and Configuration History

These capabilities are supported for System Storage DS3000, System Storage DS4000, System Storage DS5000, SAN Volume Controller, Storwize V7000, and XIV.

Tivoli Storage Productivity Center Storage Tier reports

Storage tier reports help you decide where to place frequently and infrequently accessed data for the best performance from your system. Infrequently accessed data can be placed on lower-cost storage systems and frequently accessed data can be placed on faster performing storage systems.

To create the reports, you must have Tivoli Storage Productivity Center Standard Edition or Tivoli Storage Productivity Center for Disk installed. You must also have Tivoli Common Reporting installed. These reports are supported for SAN Volume Controller and Storwize V7000. To generate the reports, you must set the back-end subsystem pool properties, then run a performance monitor job for the SAN Volume Controller or Storwize V7000. You can then display the reports using Tivoli Common Reporting.

If you have installed any IBM service offerings (for example, Storage Tiering Activity Reporter (STAR), you cannot migrate data from STAR to Tivoli Storage Productivity Center.

Tivoli Storage Productivity Center reports

For SAN Volume Controller clusters and Storwize V7000 systems, you can drill down from storage subsystems and expose managed disks, managed disk groups, and virtual disks. For Storwize V7000 systems, you can also drill down from the internal disks node to view the internal storage on the system. You can also generate a Storage Subsystem Disks report to view the internal storage on the system.

SAN Planner

The SAN Planner now supports workload profiles and replication planning for XIV systems.

The SAN Planner is able to generate and implement a storage replication plan for all three replication types on the XIV. These replication types are as follows:

- Snapshot
- Metro Mirror with Failover/Failback
- Global Mirror with Failover/Failback.

Operating system support

Tivoli Storage Productivity Center supports the following new operating systems for the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication servers:

- Red Hat Enterprise Linux Server and Advanced Platform 6 for x86-64
- ► AIX 7.1(64-bit) IBM POWER5, IBM POWER6®, IBM POWER7®

Tivoli Storage Productivity Center supports the following operating systems for the Storage Resource agents:

- ► HP-UX 11i v3 with Itanium
- ► Red Hat Enterprise Linux Advanced Platform Version 5.4
- ► AIX VIOS 2.1, 2.2
- ▶ Windows 2008 R2 (Standard Edition, Data Center Edition, and Enterprise Edition)
- Sun Solaris 9 and 10 (SPARC architecture)
- Windows 2008 Standard Edition
- ▶ Windows 2008 SP2 (Standard Edition, Data Center Edition, and Enterprise Edition)

The Tivoli Storage Productivity Center installation wizard does not support the installation of the stand-alone GUI on Windows 7. As an alternative, you can run the GUI using Java Web Start with Java 6. Java 6 provides additional features to work with the enhanced security of Windows 7.

Reference: To find the Tivoli Storage Productivity Center support matrix for various components and versions, go to the Tivoli Storage Productivity Center support website:

http://www-947.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition

See "Tivoli Storage Productivity Center support matrix" on page 792 for more details.

1.5 New functions since version Tivoli Storage Productivity Center for Replication 4.1

This section lists new functions in Tivoli Storage Productivity Center for Replication since V4.1 and is divided into three segments;

- New for Tivoli Storage Productivity Center for Replication Version 4.2.0
- ▶ New for Tivoli Storage Productivity Center for Replication Version 4.2.1
- New for Tivoli Storage Productivity Center for Replication Version 4.2.2

1.5.1 New for IBM Tivoli Storage Productivity Center for Replication 4.2.0

Tivoli Storage Productivity Center for Replication 4.2 adds the following new features, functions, and enhancements since Tivoli Storage Productivity Center V4.1. More details about Tivoli Storage Productivity Center for Replication are in Chapter 10, "Tivoli Storage Productivity Center for Replication" on page 349.

Open HyperSwap replication

Open HyperSwap replication is a special Metro Mirror replication method designed to automatically failover I/O from the primary logical devices to the secondary logical devices in the event of a primary disk storage system failure. This function can be done with minimal disruption to the applications that are using the logical devices.

Open HyperSwap replication applies to both planned and unplanned replication swaps. When a session has Open HyperSwap enabled, an I/O error on the primary site automatically causes the I/O to switch to the secondary site without any user interaction and with minimal application impact. In addition, while Open HyperSwap is enabled, the Metro Mirror session supports disaster recovery. If a write is successful on the primary site but is unable to get replicated on the secondary site, IBM Tivoli Storage Productivity Center for Replication suspends all replication for the session, thus ensuring that a consistent copy of the data exists on the secondary site. If the system fails, this data might not be the latest data, but the data will be consistent and allow the user to manually switch host servers to the secondary site.

Soft removal of hardware relationships

When you remove a copy set from IBM Tivoli Storage Productivity Center for Replication, you can choose to keep the hardware relationships on the storage systems. This way is useful when you want to migrate from one session type to another or when resolving problems.

Downloading log packages from the graphical user interface

To aid in speedy diagnostics of IBM Tivoli Storage Productivity Center for Replication anomalies, you can download a log package to the local system from the graphical user interface. You no longer need to log into the IBM Tivoli Storage Productivity Center for Replication server to collect the log package.

Global Mirror and Metro Mirror Path Manager

This feature provides peer-to-peer remote copy (PPRC) path support. You can do the following tasks with the Path Manager:

- Specify what ports to use when establishing the PPRC paths and keep that information persistent for use when the path is terminated because of a peer-to-peer suspend operation.
- Specify port pairings in a simple CSV file format to establish PPRC data paths. The specified port pairings are used whenever new paths must be established.

Additional details available for Global Mirror sessions

For Global Mirror sessions, the Session Details panel now includes the Global Mirror Info tab. This tab displays information about the Global Mirror session, including information about the Global Mirror master, consistency groups that have been formed, and data exposure time.

SAN Volume Controller session enhancements

IBM Tivoli Storage Productivity Center for Replication supports the following items:

- ► SAN Volume Controller space-efficient volumes in all IBM Tivoli Storage Productivity Center for Replication SAN Volume Controller sessions.
- ► SAN Volume Controller space-efficient volumes are intended to be used as IBM FlashCopy® targets.
- ► SAN Volume Controller incremental FlashCopy in the IBM Tivoli Storage Productivity Center for Replication FlashCopy, Metro Mirror with practice, and Global Mirror with practice sessions.

DS8000 session enhancements

IBM Tivoli Storage Productivity Center for Replication supports the following items:

- ▶ DS8000 extent space-efficient volumes on all IBM Tivoli Storage Productivity Center for Replication DS8000 sessions
- ▶ IBM Tivoli Storage Productivity Center for Replication displays whether a volume is extent space-efficient. There are certain restrictions regarding whether a space-efficient volume can be placed in a copy set. This restriction is based on the DS8000 microcode.
- ► Multiple Global Mirror sessions in a storage system so you can create multiple sessions and individually manage (start, suspend, recover, and so on) data assigned to different hosts or applications.

DB2 no longer supported as the data store for operational data

With Version 4.2, IBM Tivoli Storage Productivity Center for Replication no longer supports DB2 as the data store for its operational data. It now uses an embedded repository for its operational data. The IBM Tivoli Storage Productivity Center for Replication 4.2 installation program automatically migrates any data in an existing and operational IBM Tivoli Storage Productivity Center for Replication DB2 database to the embedded repository as part of upgrading to IBM Tivoli Storage Productivity Center for Replication 4.2 from an earlier version. New IBM Tivoli Storage Productivity Center for Replication 4.2 installations use the embedded repository by default.

1.5.2 New for IBM Tivoli Storage Productivity Center for Replication 4.2.1

This section describes the new products available.

SAN Volume Controller Version 6.1

Tivoli Storage Productivity Center for Replication supports SAN Volume Controller Version 6.1 and IBM Storwize V7000 as described in Tivoli Storage Productivity Center.

IBM Storwize V7000

Tivoli Storage Productivity Center for Replication supports SAN Volume Controller Version 6.1 and IBM Storwize V7000 as described in Tivoli Storage Productivity Center.

1.5.3 New for IBM Tivoli Storage Productivity Center for Replication 4.2.2

This section describes the new products available.

Tivoli Storage Productivity Center for Replication 4.2.2 supports IBM XIV Storage System

You can now use the following session types for an XIV system.

Snapshot

Snapshot is a new session type that creates a point-in-time copy (snapshot) of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

Metro Mirror Failover/Failback

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

Global Mirror Failover/Failback

Global Mirror is a method of asynchronous, remote data replication that operates between two sites that are over 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

Tivoli Storage Productivity Center installation on Windows

In this chapter, we show the installation steps of the Tivoli Storage Productivity Center V4.2 on the Windows platform. Of the available installation paths, Typical and Custom, we describe the Custom installation in our environment. We also list the Tivoli Storage Productivity Center for Replication installation considerations.

This installation process documented in this book is based on an environment where clean servers are available for the installation.

2.1 Tivoli Storage Productivity Center installation

Tivoli Storage Productivity Center has an installation wizard that guides you through the installation of the Tivoli Storage Productivity Center servers and agents. The installation described in this chapter is not specific to a Tivoli Storage Productivity Center license. All editions use the same code base and as such all the panels look alike.

The prerequisite components have to be installed prior to invoking the installation wizard.

2.1.1 Installation overview

To get Tivoli Storage Productivity Center V4.2 to work, use the following steps:

- Check that the system meets the prerequisites.
 See 2.2, "Preinstallation steps for Windows" on page 24.
- Install the prerequisite components.
 See 2.3, "Installing the prerequisite for Windows" on page 28.
- 3. Install Tivoli Storage Productivity Center components.

 See 2.4, "Installing Tivoli Storage Productivity Center components" on page 38.
- 4. Install Tivoli Storage Productivity Center agents. See 2.4.3, "Agent installation" on page 66.

You need to understand the difference between an agent installation and a deployment:

- ▶ We use the term *installation* if the agent is locally installed with a GUI or CLI installer
- ▶ We use the term *deployed* when the Tivoli Storage Productivity Center server is running and pushes the agent onto a server, without being locally logged in on that system.

You can install all the Tivoli Storage Productivity Center components using Typical installation or Custom installation.

Typical installation

With the Typical installation, you install all the components of the Tivoli Storage Productivity Center on the local server in one step, although you can still decide which components to install:

- Server: Data server, Device server, Replication Manager, and Tivoli Integrated Portal
- ► Clients: Tivoli Storage Productivity Center GUI
- ► Storage Resource agent

The "Typical" installation path is ideal for small and medium sized environments where you do not need to customize installation details and the defaults are sufficient.

The "Custom" path is preferred for large environments, or when you need to have greater control over installation choices, such as these:

- Size and location of the Tivoli Storage Productivity Center database, tables, logs, and so on
- ► Greater flexibility with user accounts and passwords for components and services

Custom installation

With the Custom installation, you install parts of Tivoli Storage Productivity Center separately; The Custom installation method provides options so you can change default settings, such as user IDs and directories. This is the installation method that we prefer.

Installing Tivoli Storage Productivity Center has the following installable components:

- Database Schema
- ► Data server and Device server
- ► Graphical user interface (GUI)
- ► Command-line interface (CLI)
- ► Storage Resource agent

Tivoli Storage Productivity Center for Replication install considerations

When the installation is about 75% complete, the installer launches the Tivoli Storage Productivity Center for Replication installation wizard, which presents you with the option to change various installation parameters. You basically have to step through the wizard, and then click **Finish** to start the Tivoli Storage Productivity Center for Replication installation process. After the process completes, click **Finish** to return to the Tivoli Storage Productivity Center installer to complete the remaining installation steps.

Installation timing

The approximate time to install Tivoli Storage Productivity Center, including Tivoli Integrated Portal, is about 60 minutes. The approximate time to install Tivoli Storage Productivity Center for Replication is about 20 minutes.

2.1.2 Product code media layout and components

In this section, we describe the contents of the product media at the time of writing. The media content differs depending on whether you are using the web images or the physical media included with the Tivoli Storage Productivity Center V4.2 package.

Passport Advantage and web media content

The web media consists of a disk image and an SRA ZIP file. The disk image is divided into four parts:

- ▶ Disk1 part 1: Contains the following Tivoli Storage Productivity Center components:
 - Database Schema
 - Data server
 - Device server
 - GUI
 - CLI
 - Storage Resource agent
- ▶ Disk1 part 2: Contains the following Tivoli Storage Productivity Center components:
 - IBM Tivoli Integrated Portal
 - IBM Tivoli Storage Productivity Center for Replication
- ▶ Disk1 part 3: Contains the following fix pack:
 - IBM Tivoli Integrated Portal Fixpack

Attention: Part 1, part 2, and part 3 are required for every Tivoli Storage Productivity Center installation and must be downloaded and extracted to a single directory.

- ▶ Disk1 part 4: Contains an optional component:
 - IBM Tivoli Storage Productivity Center Monitoring Agent for IBM Tivoli Monitoring

Important: On Windows, ensure that the directory name where the installation images reside has no spaces or special characters, otherwise, the Tivoli Storage Productivity Center installation will fail. This issue can happen if you have a directory similar to the following name, for example:

C:\tpc 42 standard edition\disk1

The SRA ZIP file contains Tivoli Storage Productivity Center Storage Resource agent. It does not include a GUI installer.

To understand how this installation method works see Chapter 8, "Storage Resource Agent" on page 263.

The content of this disk is as follows:

- Directory: readme
- Directory: sra
- ▶ File: version.txt

In addition to the images we mentioned, the media contains the following images:

- Tivoli Storage Productivity Center Storage National Language Support
- ► IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity License, which is available for Windows, Linux and AIX
- ► IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity License, which is available for Windows, Linux and AIX

Physical media

The physical media included with the Tivoli Storage Productivity Center V4.2 product consists of a DVD and a CD. The DVD contains the Disk1 part 1 and Disk1 part 2 content described in "Passport Advantage and web media content" on page 23. The physical media CD is the same as the web Disk2 media.

2.2 Preinstallation steps for Windows

Before you install the Tivoli Storage Productivity Center V4.2, note the following prerequisite information:

- ► For Tivoli Storage Productivity Center V4.2, IBM DB2 Universal Database[™] (UDB), Enterprise Server Edition is the only prerequisite component with the following supported levels:
 - V9.1 (Fix Pack 2 or later, not supported on Windows 2008 R2; use DB2 9.5 or 9.7)
 - V9.5 (Fix Pack 6, not supported)
 - V9.7 (Fix Pack 1 to 3a, not supported)

For the most current Tivoli Storage Productivity Center 4.2 platform support, see this website:

https://www-304.ibm.com/support/docview.wss?uid=swg27019380

► Starting from Tivoli Storage Productivity Center V4.1, the installation of Tivoli Agent Manager is optional. You are required to install it only if you need to use Data or Fabric agents on platforms that are not supported with Storage Resource agents.

Agent Manager 1.3.2 (any sub-version) supports DB 9.1. For DB2 9.5 support, you need to use Agent Manager Version 1.3.2.30, which is included with Tivoli Storage Productivity Center 4.1.1. If you are planning to use DB2 9.7, you must install Agent Manager 1.4.x or later, which is included with Tivoli Storage Productivity Center Version 4.2.x.

Order of prerequisite component installation

Follow this order when installing the prerequisite components:

- DB2 UDB
- 2. Optional: Tivoli Agent Manager, but this can be also installed later when Tivoli Storage Productivity Center is already running

2.2.1 Verifying system hardware and software prerequisites

For hardware and software prerequisites, see the Tivoli Storage Productivity Center support site.

http://www-947.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition

2.2.2 Verifying primary domain name systems

Before you start the installation, verify whether a primary domain name system (DNS) suffix is set. This determination can require restarting the computer.

Verify the primary DNS name as follows:

- 1. Right-click My Computer on your desktop.
- 2. Click Properties.

The System Properties panel is displayed (Figure 2-1).

3. Click the Computer Name tab. On the panel that is displayed, click Change.

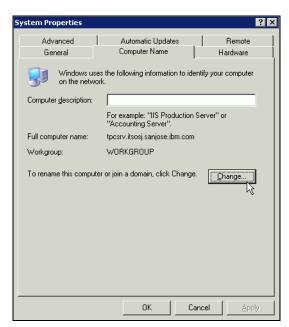


Figure 2-1 System Properties

4. Enter the host name in the Computer name field. Click **More** to continue (Figure 2-2).

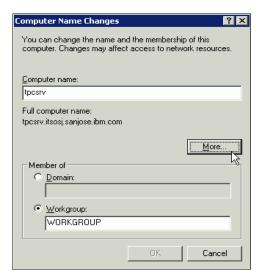


Figure 2-2 Computer name

5. In the next panel, verify that Primary DNS suffix field shows the correct domain name. Click **OK** (Figure 2-3).

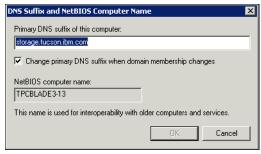


Figure 2-3 DNS domain name

6. If you made any changes, you must restart your computer for the changes to take effect (Figure 2-4).

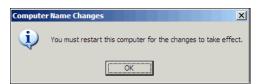


Figure 2-4 Restart the computer for changes to take effect

2.2.3 Activating NetBIOS settings

If NetBIOS is not enabled on Microsoft Windows 2003, then GUID is not generated. You must verify and activate NetBIOS settings.

Complete the following steps:

 On your Tivoli Storage Productivity Center Server, go to Start → Control Panel → Network Connections.

- 2. Select your Local Area Connections.
- 3. From the Local Area Connection Properties panel, double-click **Internet Protocol** (TCP/IP). The next panel is Internet Protocol (TCP/IP) Properties.
- 4. Click **Advanced** as shown in Figure 2-5.

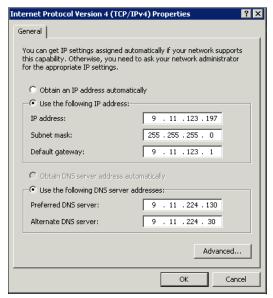


Figure 2-5 TCP/IP properties

5. On the WINS tab, select Enable NetBIOS over TCP/IP and click OK (Figure 2-6).

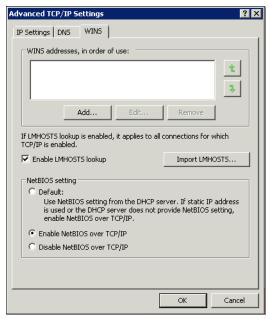


Figure 2-6 Advanced TCP/IP properties

2.2.4 User IDs and passwords to be used and defined

For considerations and information about the user IDs and passwords defined or set up during Tivoli Storage Productivity Center Productivity Center installation, see the planning chapter in the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide*, SC27-2337. We have added a table in Appendix B, "Worksheets" on page 803 that can help you track the created user ID and storage subsystem information.

Tips:

- A good practice is to use the worksheets in Appendix B, "Worksheets" on page 803 to record the user IDs and passwords used during the installation of Tivoli Storage Productivity Center.
- ► If you want to use another user ID, create it before beginning the installation and ensure that it has administrator rights.

2.3 Installing the prerequisite for Windows

In this section, we show how to install the Tivoli Storage Productivity Center prerequisites on Windows. We perform a typical installation of DB2 Enterprise Server Edition Version 9.7.

Before beginning the installation, log on to your system as a local administrator with Administrator authority.

2.3.1 DB2 installation

Install DB2 as follows:

1. Insert the IBM DB2 Installer CD into the CD-ROM drive.

If Windows autorun is enabled, the installation program starts automatically. However, if it does not start automatically, open Windows Explorer, go to the DB2 Installation image path, and double-click the **setup.exe** file.

Attention: Only the user ID that installed the DB2 product has the privilege to issue the **db2start** and **db2stop** commands.

On the Welcome panel, shown in Figure 2-7, select **Install a Product** to proceed with the installation.

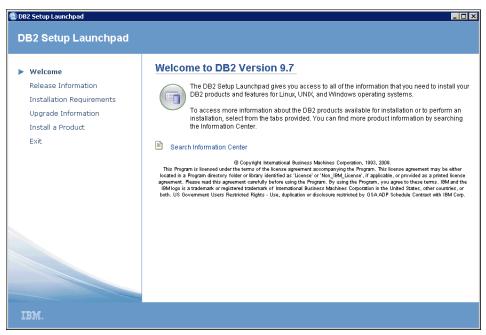


Figure 2-7 DB2 Setup Welcome panel

2. At the next panel, select the DB2 product to be installed. Select **DB2 Enterprise Server Edition Version 9.7**, and click **Install New**, as shown in Figure 2-8.



Figure 2-8 Select product

3. The DB2 Setup wizard panel opens, as shown in Figure 2-9. Click Next to proceed.

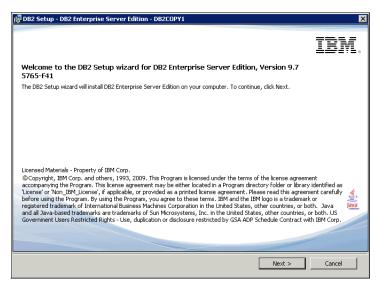


Figure 2-9 Setup wizard

4. Read the software license agreement and, if you agree, select I accept the terms in the license agreement (Figure 2-10), and click Next.

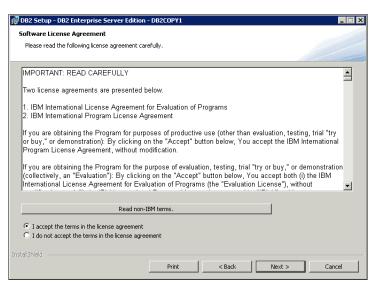


Figure 2-10 License agreement

5. To select the installation type, accept the default of **Typical** and click **Next** to continue (Figure 2-11).

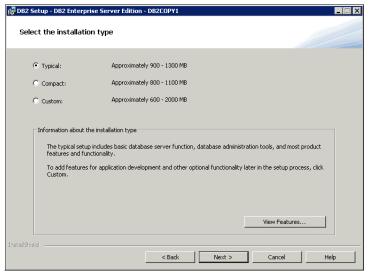


Figure 2-11 Typical installation

6. Select Install DB2 Enterprise Server Edition on this computer and save my settings in a response file (see Figure 2-12). Specify the path and the file name for the response file in the Response file name field. The response file will be generated at the end of the installation process and it can be used to perform additional silent installations of DB2, using the same parameters specified during this installation. Click Next to continue.

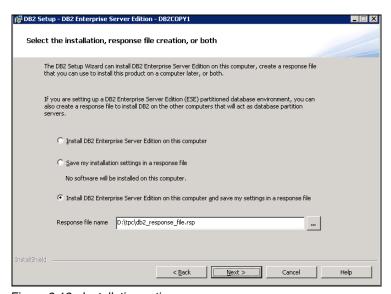


Figure 2-12 Installation action

7. The panel in Figure 2-13 shows the default values for the drive and directory to be used as the installation folder. You can change these or accept the defaults, then click **Next** to continue. In our installation, we accept to install on the C: drive.

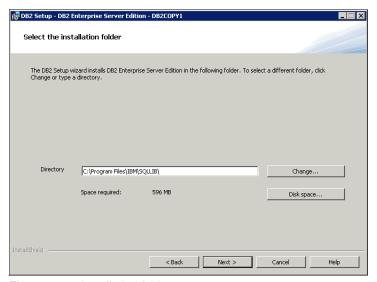


Figure 2-13 Installation folder

 The next panel requires user information for the DB2 Administration Server; it can be a Windows domain user. If it is a local user, select **None - use local user account** for the Domain field.

The user name field is prefilled with a default user name. You can change it or leave the default and type the password of the DB2 user account that you want to create (Figure 2-14). Keep the check mark in the **Use the same user name and password for the remaining DB2 services** check box, and click **Next** to continue.

DB2 creates a user with the following administrative rights:

- Act as a part of an operating system.
- Create a token object.
- Increase quotas.
- Replace a process-level token.
- Log on as a service.

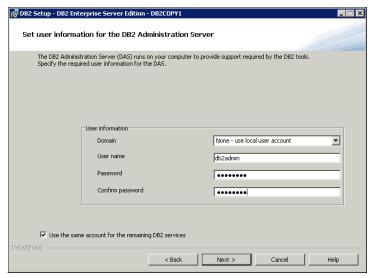


Figure 2-14 User Information

9. In the Configure DB2 instances panel, accept the default and click **Next** to continue (Figure 2-15).

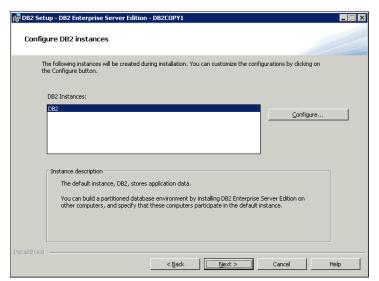


Figure 2-15 Configure DB2 instances

10.On the next panel, which allows you to specify options to prepare the DB2 tools catalog, accept the defaults, as shown in Figure 2-16. Verify that the Prepare the DB2 tools catalog check box on this computer is not selected. Click Next to continue.

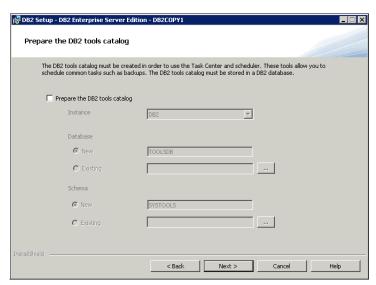


Figure 2-16 Prepare db2 tools catalog

11.On the next panel, shown in Figure 2-17, you can set the DB2 server to send notifications when the database needs attention. Ensure that the Set up your DB2 server to send notification check box is not selected, and then click Next to continue.

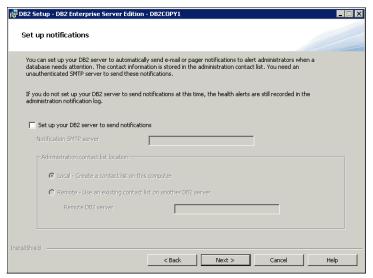


Figure 2-17 Health Monitor

12. Accept the defaults for the DB2 administrators group and DB2 users group in the "Enable operating system security for DB2 objects" panel shown in Figure 2-18 and click **Next** to proceed.

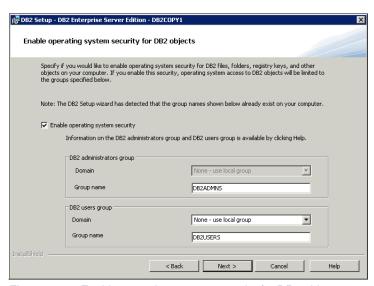


Figure 2-18 Enable operating system security for DB2 objects

13. Figure 2-19 shows a summary of what will be installed, based on your input. Review the settings and click **Install** to continue.

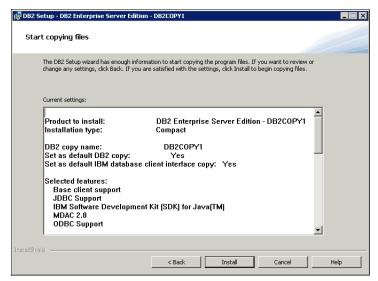


Figure 2-19 Summary panel

As the DB2 installation proceeds, a progress status bar is displayed (Figure 2-20).

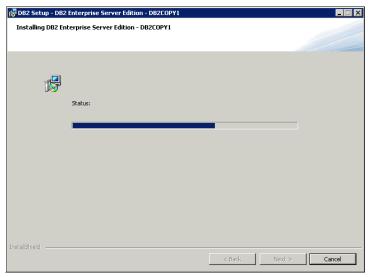


Figure 2-20 DB2 Enterprise Server Edition installation progress

14. When the setup completes, click **Next**, as shown in Figure 2-21.

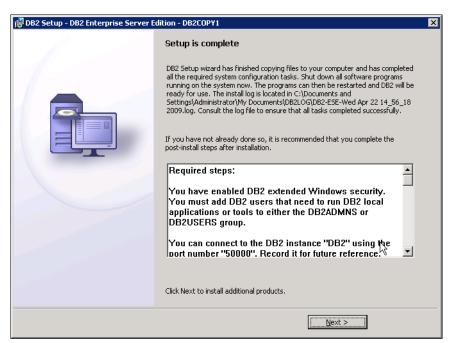


Figure 2-21 DB2 setup summary panel

15. At the next panel (Figure 2-22), you can choose to install additional products. In our installation, we clicked **Finish** to exit the DB2 setup wizard.



Figure 2-22 DB2 setup final panel

16.Click Exit on the DB2 Setup Launchpad (Figure 2-23) to complete the installation.

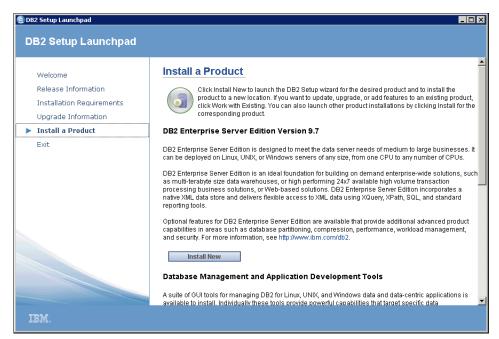


Figure 2-23 DB2 Setup Launchpad

Tip: When applying a DB2 fix pack, you must select the DB2 Enterprise Server Edition and not other options, as depicted in Figure 2-23.

2.3.2 Verifying the installation

Verify the DB2 installation as follows:

Launch a DB2 Command window by selecting Start → IBM DB2 → DB2COPY1 (Default) → Command Line Tools → Command Window (Figure 2-24).

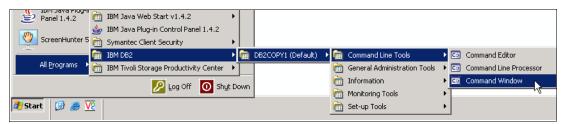


Figure 2-24 DB2 Command Windows

Create the SAMPLE database by entering the db2samp1 command, as shown in Figure 2-25.

```
C:\Program Files\IBM\SQLLIB\BIN>db2samp1

Creating database "SAMPLE"...
Existing "SAMPLE" database found...
The "-force" option was not specified...
Attempt to create the database "SAMPLE" failed.
SQL1805N The database alias "SAMPLE" already exists in either the local database directory or system database directory.

'db2samp1' processing complete.

C:\Program Files\IBM\SQLLIB\BIN>_
```

Figure 2-25 Create the SAMPLE database

3. Enter the following DB2 commands, which connect to the SAMPLE database, issue a simple SQL query, and reset the database connection:

```
db2 connect to sample
db2 "select * from staff where dept = 20"
db2 connect reset
```

The result of these commands is shown in Figure 2-26.

```
:\Program Files\IBM\SQLLIB\BIN>db2 connect to sample
   Database Connection Information
                         = DB2/NT 9.7.0
 Database server
SQL authorization ID
Local database alias
                         = ADMINIST...
C:\Program Files\IBM\SQLLIB\BIN>db2 "select * from staff where dept = 20"
D
                         JOB YEARS SALARY
       NAME
                  DEPT
                                                  COMM
                                        98357.50
    10 Sanders
                      20 Mgr
                      20 Sales
20 Clerk
                                        78171.25
43504.60
    20 Pernal
                                     8
                                                      612.45
    80 James
                                                      128.20
   190 Sneider
  4 record(s) selected.
C:\Program_Files\IBM\SQLLIB\BIN>db2 connect reset
          The SQL command completed successfully.
C:\Program Files\IBM\SQLLIB\BIN>
```

Figure 2-26 DB2 command results

2.4 Installing Tivoli Storage Productivity Center components

Now that all the prerequisites are installed, we can install the Tivoli Storage Productivity Center components, keeping in mind that with Tivoli Storage Productivity Center V4.2, both Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication are installed. We divide the installation in two steps:

- Create the Database Schema.
- 2. Install the Data server and the Device server.

The reason for two stages is because if you install all the components in one step and any part of the installation fails for any reason (for example, space or passwords), the installation suspends and rolls back, uninstalling all the previously installed components. Other than that, you could also install the schema and Tivoli Storage Productivity Center at the same time.

2.4.1 Creating the Database Schema

Before starting the installation, verify that a supported version of DB2 Enterprise Server Edition has been installed and started.

Important: Log on to your system as a local administrator with database authority.

Complete the following steps:

- 1. If Windows autorun is enabled, the installation program starts automatically. If it does not start automatically, open Windows Explorer and go to the Tivoli Storage Productivity Center CD–ROM drive or directory. Double-click the setup.exe file.
- 2. Select a language and click **OK** (Figure 2-27).



Figure 2-27 Language selection panel

3. Read the international program license agreement and, if you agree, select I accept the terms in the license agreement, and then click Next to continue (see Figure 2-28).

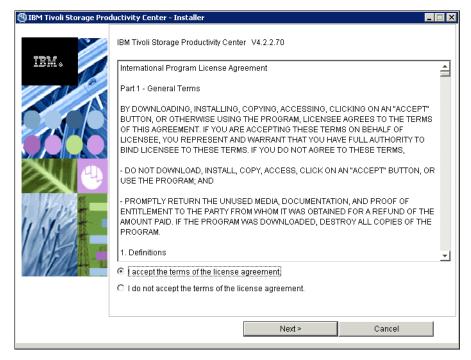


Figure 2-28 License agreement

- 4. Figure 2-29 shows how to select typical or custom installation. You have the following options:
 - Typical installation:

With this option, you install all components on the same computer by selecting **Servers, Agents,** and **Clients**.

Custom installation:

With this option, install the database schema, the Tivoli Storage Productivity Center server, CLI, GUI and Storage Resource agent separately.

Installation licenses:

This option installs the Tivoli Storage Productivity Center licenses. The Tivoli Storage Productivity Center license is on the CD. You only need to run this option when you add a license to a Tivoli Storage Productivity Center package that has already been installed on your system.

For example, if you have installed Tivoli Storage Productivity Center for Data package, the license will be installed automatically when you install the product. If you decide to later enable Tivoli Storage Productivity Center for Disk, run the installer and select Installation licenses. This option will allow you to install the license key from the CD. You do not have to install the Tivoli Storage Productivity Center for Disk product.

In this chapter, we describe *Custom installation*. Select the directory where you want to install Tivoli Storage Productivity Center. A default installation directory is suggested; you can accept it or change it, and then click **Next** to continue.

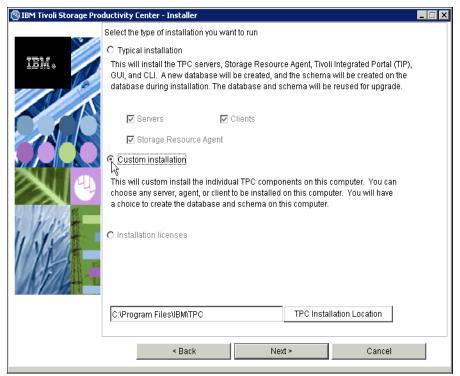


Figure 2-29 Custom installation selected

5. In the Custom installation, you can select all the components in the panel shown in Figure 2-30. By default, all components are checked. In our scenario, we show the installation in stages, so we only select the option to **Create database schema**, and click **Next** to proceed.

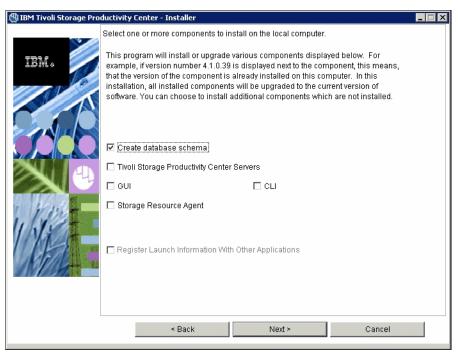


Figure 2-30 Custom installation component selection

To start the database creation, specify a DB2 Database administrator user ID and password. We suggest that you use the same DB2 user ID that you created when you installed DB2. Click Next. See Figure 2-31.

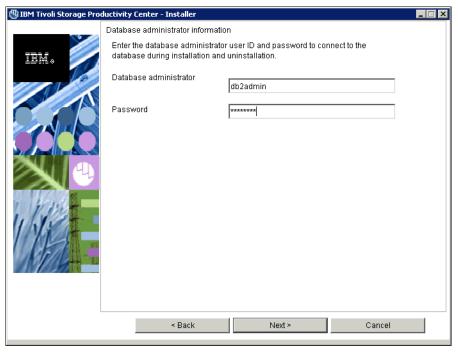


Figure 2-31 DB2 user and password

Important: The Database administrator user ID that you provide should be part of the DB2ADMNS group, because only those users are allowed to perform the actions that are required to create a new database and install the schema into that database.

If the user ID you enter is not part of the DB2ADMNS group, the installation will likely fail at about 7% completion (Figure 2-32).

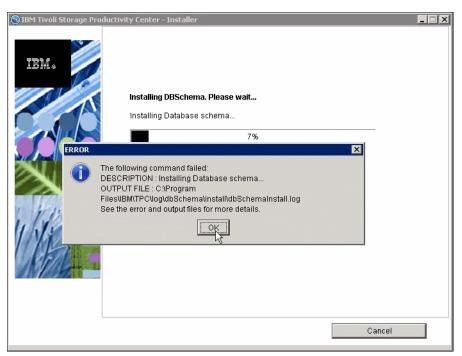


Figure 2-32 DB2 error during schema installation

7. Enter your DB2 user ID and password again. This ID does *not* have to be the same ID as the first one. Make sure that the **Create local database** option is selected. See Figure 2-33. By default, a database named TPCDB is created. Click **Database creation details** to continue.

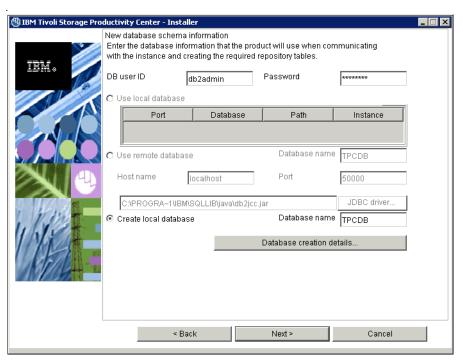


Figure 2-33 DB2 user and create local database

Use the panel shown in Figure 2-34 to change the default space assigned to the database. Review the defaults and make any changes. In our installation, we accepted the defaults.

For better performance, use the following considerations:

- Allocate TEMP DB on a separate physical disk from the Tivoli Storage Productivity Center components.
- Create larger Key and Big Databases.

Select **System managed (SMS)** and click **OK** and then **Next** to proceed. To understand the advantage of an SMS database versus a DMS database or the Automatic Storage, see "Selecting an SMS or DMS table space" on page 800.

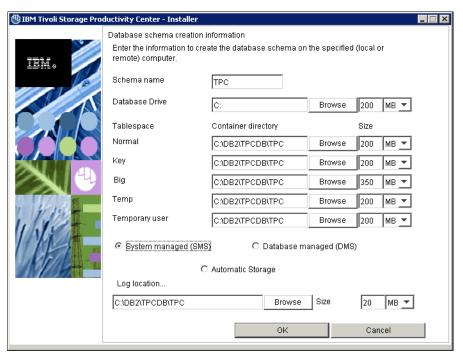


Figure 2-34 DB schema space

Character length: The Tivoli Storage Productivity Center schema name cannot be longer than eight characters.

8. The Tivoli Storage Productivity Center installation information that you selected is listed, as shown in Figure 2-35. Click **Install** to continue.

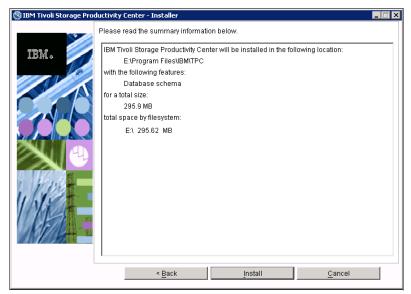


Figure 2-35 Tivoli Storage Productivity Center installation information

Figure 2-36 shows the Database Schema installation progress panel. Wait for the installation to complete.

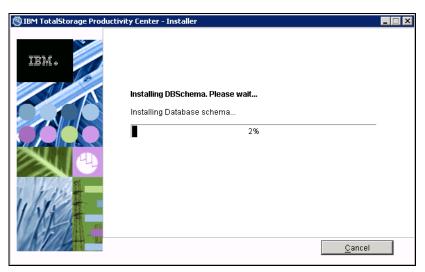


Figure 2-36 installing database schema

9. At the successfully installed message (Figure 2-37), click **Finish** to continue.

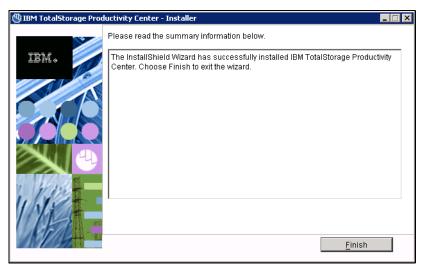


Figure 2-37 Installation summary information

Verifying the installation

To check the installation, choose $Start o All\ Programs o IBM\ DB2 o DB2COPY1$ (Default) o General Administration Tools o Control Center, to start the DB2 Control Center. Under All Databases, verify that you have at least a database named TPCDB, as shown in Figure 2-38.

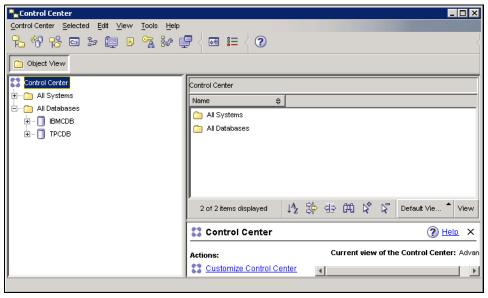


Figure 2-38 Verifying DB2 installation

Important: Do not edit or modify anything in the DB2 Control Center. Doing so can cause serious damage to your table space. Simply use the DB2 Control Center to browse your configuration.

Log files

Check for errors and Java exceptions in the log files at the following locations:

- <InstallLocation>\TPC.log
- <InstallLocation>\log\dbSchema\install

For Windows, the default installation location is as follows:

c:\Program Files\IBM\TPC

At the end of the log files, verify that a "success" message indicates a successful installation.

2.4.2 Installing Tivoli Storage Productivity Center components

In this step, we perform a custom installation to install the following components:

- Data server
- Device server
- ► GUI
- ► CLI

Important: Do not install the Storage Resource agent at this time. Installing any Storage Resource agent by the installer requires that you also uninstall the Storage Resource agent using the installer. In most cases, using the Tivoli Storage Productivity Center GUI to deploy agents (instead of installing them) is the more flexible approach.

During this process, two additional components are also installed:

- Tivoli Integrated Portal
- ► Tivoli Storage Productivity Center for Replication

Preinstallation tasks

To install Data server and Device server components, you must log on to the Windows system with a user ID that has the following rights, which any user that is part of the DB2ADMNS group has automatically:

- ► Log on as a service.
- Act as part of the operating system.
- Adjust memory quotas for a process.
- Create a token object.
- ► Debug programs.
- Replace a process-level token.

Be certain that the following tasks are completed:

- The Database Schema must be installed successfully to start the Data server installation.
- ► The Data server must be successfully installed prior to installing the GUI.
- ► The Device server must be successfully installed prior to installing the CLI.

Local Database Validation error

After a successful DB2 9.7 installation on 64-bit Windows 2003 and Windows 2008 servers, the database instance will not be recognized by the Tivoli Storage Productivity Center installer without a reboot of the server. During Tivoli Storage Productivity Center V4.2 installation or upgrade, a message window might open that contains the following message:

Local database validation: No valid local database found on the system for installation of the Data Server, Device server or creation of the database schema.

Information about this message is provided at the following location:

https://www-304.ibm.com/support/docview.wss?uid=swg21452614

Custom installation

To perform a custom installation, complete the following steps:

- 1. Start the Tivoli Storage Productivity Center installer.
- 2. Choose the language to be used for installation.
- 3. If you agree, accept the terms of the License Agreement.
- 4. Select the Custom installation.
- 5. Select the components you want to install. In our scenario, we select the Servers, GUI, CLI as shown in Figure 2-39. Notice that the **Create database schema** check box is *not* selected. Click **Next** to continue.



Figure 2-39 Installation selection

6. If you are running the installation on a system with at least 4 GB but less than 8 GB of RAM, a warning message is issued (Figure 2-40). Click **OK** to close the message and continue.



Figure 2-40 Memory warning message

7. In the Database administrator information panel (Figure 2-41), the DB2 user ID and password are filled in because we used them to create the Database Schema. Click **Next**.

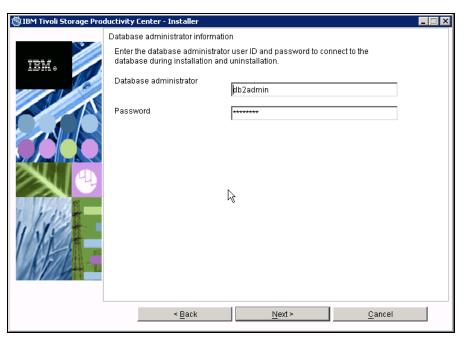


Figure 2-41 DB2 User ID and password

The user ID is saved to the install/uninstall configuration files, therefore, if the password changes from the time you first installed a Tivoli Storage Productivity Center component, an incorrect password might be populated into this panel.

8. We want to use the database TPCDB that we created in the previous section on the same server. Therefore, we select **Use local database** and click **Next** to continue (Figure 2-42).

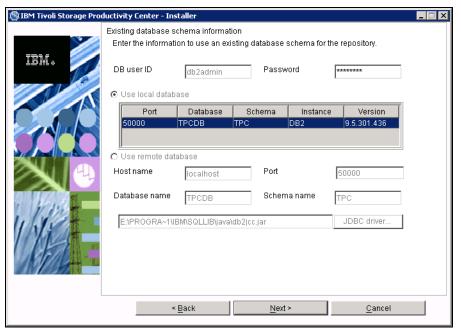


Figure 2-42 Use local database selection

Tivoli Storage Productivity Center can also run having the DB schema installed on another server. In this case, you have to install the Tivoli Storage Productivity Center schema on that server following the procedure documented in 2.4.1, "Creating the Database Schema" on page 39. Then, installing the other Tivoli Storage Productivity Center components, you must select the **Use remote database** option and specify the host name of the server running the DB2 Manager. The other fields must be prefilled, as shown in Figure 2-43. Verify their values and click **Next**.

Tip: If you have the Tivoli Storage Productivity Center schema already installed locally, the option of using a remote database is disabled. You have to uninstall the local copy and rerun the installation program to enable the remote database option.

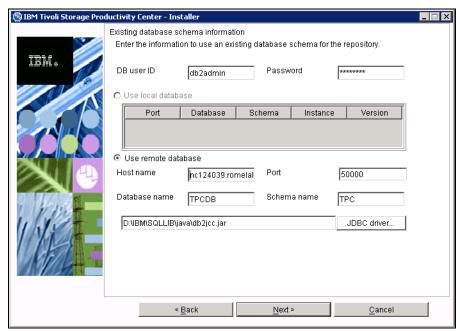


Figure 2-43 Remote database selection

If you selected to use a remote database, a warning message is issued (Figure 2-44), reminding you to ensure that the remote DB2 instance is running before proceeding.

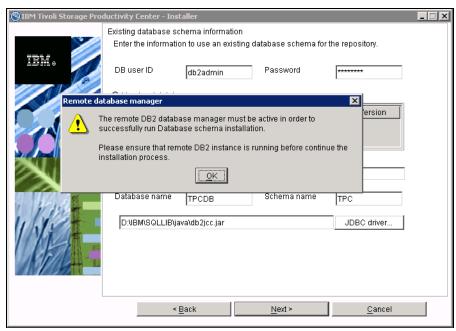


Figure 2-44 Ensure that DB2 is running on the remote system

- In the panel shown in Figure 2-45, enter the following information, and click Next to continue:
 - Data Server Name:

Enter the fully qualified host name of the Data server.

Data Server Port:

Enter the Data server port. The default is 9549.

- Device Server Name:

Enter the fully qualified host name of the Device server.

Device Server Port:

Enter the Device server port. The default is 9550.

- TPC Superuser:

Enter the name of a OS group that will be granted the superuser role within Tivoli Storage Productivity Center.

Override: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the value that you enter for LDAP Tivoli Storage Productivity Center Administrator group overrides the value that you entered here for the TPC superuser. For helpful LDAP information see Appendix D, "LDAP considerations" on page 833.

Host Authentication Password:

This password is used for internal communication between Tivoli Storage Productivity Center components, such as the Data server and the Device server.

This password can be changed when you right-click **Administrative Services** → **Services** → **Device Server** → **Device Server**, and select **Change Password**.

Data Server Account Password:

For Windows only. Tivoli Storage Productivity Center installer will create an ID called *TSRMsrv1* with the password you specified here to run the Data server service. The display name for the Data server in Windows Services panel is as follows:

IBM Tivoli Storage Productivity Center - Data Server

IBM WebSphere® Application Server admin ID and password:

This user ID and password is required by the Device server to communicate with the embedded WebSphere.

You can use the same user as the one that was entered on the panel shown in Figure 2-42 on page 49.

Override: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the value entered for the LDAP Tivoli Storage Productivity Center Administrator group overrides the value you entered here for the WebSphere Application Server admin ID and password.

- If you click Security roles, the Advanced security roles mapping panel opens. You can assign a Windows OS group to a role group for each Tivoli Storage Productivity Center role that you want to make an association with, so you can have separate authority IDs to do various Tivoli Storage Productivity Center operations. The operating group must exist before you can associate a Tivoli Storage Productivity Center role with a group. You do not have to assign security roles at installation time; you can assign these roles after you have installed Tivoli Storage Productivity Center.
- If you click NAS discovery, the NAS discovery information panel opens. You can enter
 the NAS filer login default user name and password and the SNMP communities to be
 used for NAS discovery. You do not have to assign the NAS discovery information at
 installation time, you can configure it after you installed Tivoli Storage Productivity
 Center.

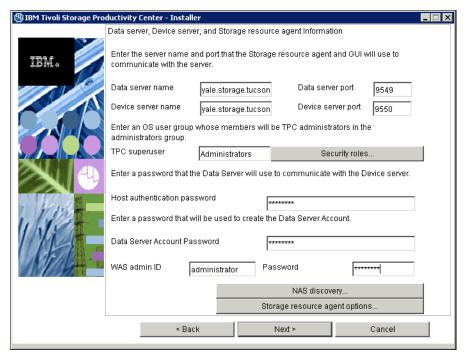


Figure 2-45 Component information for installation

- 10. In the next panel (Figure 2-46), you can select an existing Tivoli Integrated Portal to use or install a new one. Because we are installing a new instance, we have to specify the installation directory and the port number. Tivoli Integrated Portal will use 10 port numbers starting from the one specified in the Port field (called *Base Port*):
 - base port
 - base port+1
 - base port+2
 - base port+3
 - base port+5
 - base port+6
 - base port+8
 - base port+10
 - base port+12
 - base port+13

The TIP Administrator ID and Password fields are prefilled with the WebSphere admin ID and password that are specified during the Device server installation (Figure 2-45 on page 52).

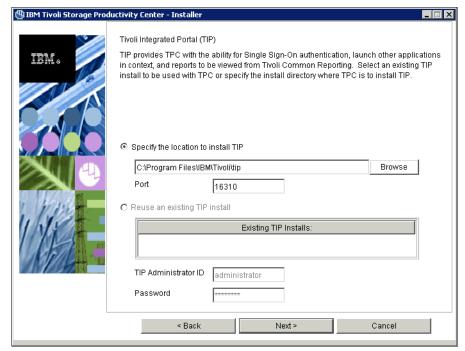


Figure 2-46 Tivoli Integrated Portal panel

- 11.On the next panel, shown in Figure 2-47, you can choose the authentication method that Tivoli Storage Productivity Center will use to authenticate the users:
 - If you want to authenticate the users against the operating system, select this option and click Next.
 - If you want to use an LDAP or Active Directory, you must have an LDAP server already installed and configured. If you decide to use this option, select the LDAP/Active directory radio button, click Next; addition panels open. For helpful LDAP information, See Appendix D, "LDAP considerations" on page 833.



Figure 2-47 Authentication type Panel

If you selected the LDAP/Active Directory option, use the following steps:

i. The panel shown in Figure 2-48 is displayed. Insert the LDAP Server host name and change the LDAP Port Number if it is not corresponding to the proposed default value. You must also fill in the Bind Distinguished Name and the Bind Password only if the anonymous binds are disabled on your LDAP server. Click **Next** to continue.

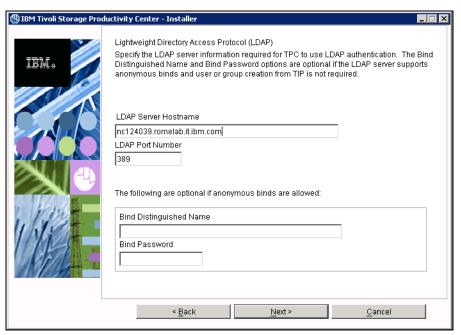


Figure 2-48 LDAP Server panel

ii. In the next panel, (Figure 2-49), you are required to insert the LDAP IBM RDN® for users and groups and the attributes that must be used to search the directory. When you click **Next**, the Tivoli Storage Productivity Center installation attempts to connect to the LDAP server to validate the provided parameters. If the validation is successful, you are prompted with the next panel; otherwise a message is issued to explain the problem.

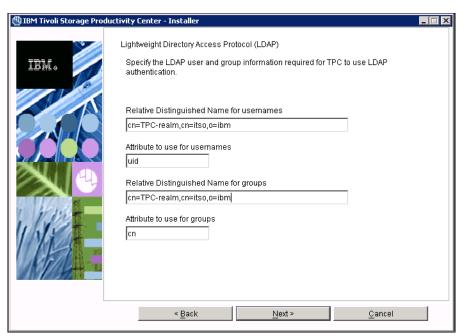


Figure 2-49 LDAP RDN details

iii. In the next panel (Figure 2-50), specify the LDAP user ID and password corresponding to the Tivoli Storage Productivity Center Administrator and the LDAP group that will be mapped to the Tivoli Storage Productivity Center Administrator group. Also in this panel, after filling in the fields and clicking **Next**, the installation program connects to the LDAP server to verify the correctness of the provided values. If the validation is successful, the next installation panel opens.

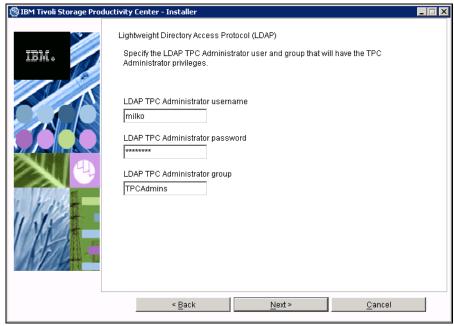


Figure 2-50 LDAP user and group for Tivoli Storage Productivity Center administration

Warning: Because of the WebSphere Application Server APAR PK77578, the LDAP Tivoli Storage Productivity Center Administrator user name value must not contain a space in it.

12. Review the summary information (Figure 2-51) and then click **Install** to continue.

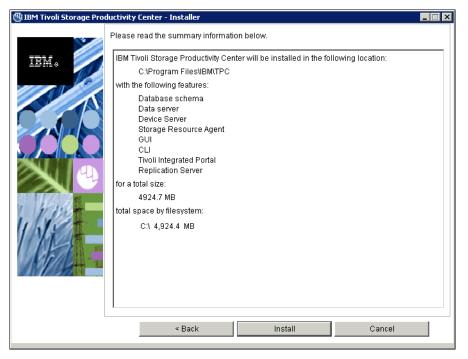


Figure 2-51 Summary of installation

The installation starts. You might see several messages:

- Messages related to Data server installation (Figure 2-52)
- Messages related to the Device server installation (Figure 2-53)
- Messages related to the Tivoli Integrated Portal installation (Figure 2-54 on page 58)

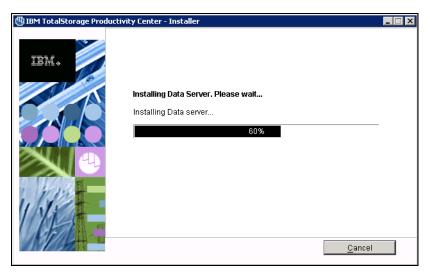


Figure 2-52 Installing Data server

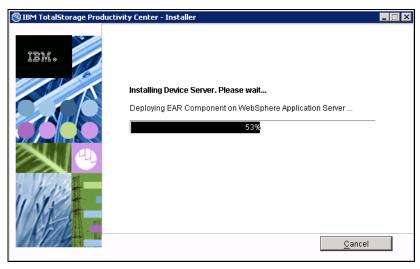


Figure 2-53 Installing Device server

Timing: Installation of the Tivoli Integrated Portal can be a time-consuming process, requiring more time than the other Tivoli Storage Productivity Center components. Installation of Tivoli Integrated Portal is completed when the process bar reaches 74%.

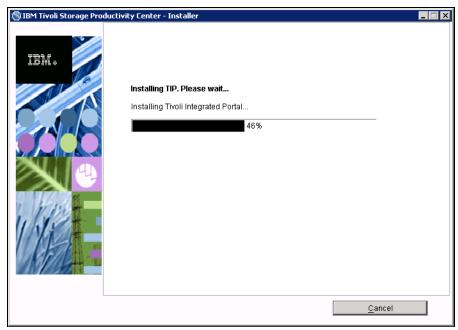


Figure 2-54 Installing Tivoli Integrated Portal

The *Tivoli Storage Productivity Center for Replication* is installed next.

Tivoli Storage Productivity Center for Replication installation

Upon completion of the Tivoli Integrated Portal installation, the Tivoli Storage Productivity Center for Replication installation is launched. The Tivoli Storage Productivity Center installation is temporarily suspended, and the panel in Figure 2-55 remains in the background while the Tivoli Storage Productivity Center for Replication installation starts (Figure 2-56.)

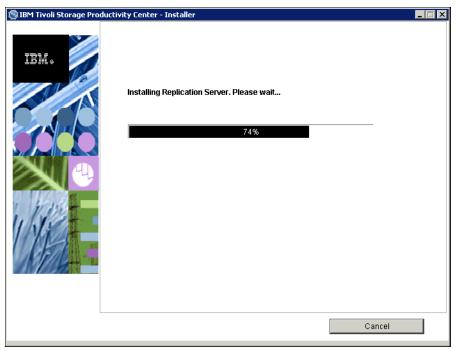


Figure 2-55 Installation panel launching the Tivoli Storage Productivity Center for Replication

Complete the following steps:

1. On the Welcome panel (Figure 2-56), click Next to proceed.



Figure 2-56 Tivoli Storage Productivity Center for replication Welcome panel

Warning: If you do not plan to use Tivoli Storage Productivity Center for Replication, do not interrupt the installation by clicking **Cancel**. Doing so results in an interruption of the installation process with a subsequent complete Tivoli Storage Productivity Center installation rollback. Complete the installation and then disable Tivoli Storage Productivity Center for Replication.

The installation wizard checks on the system prerequisites to verify that the operating system is supported and the appropriate fix packs are installed (Figure 2-57).

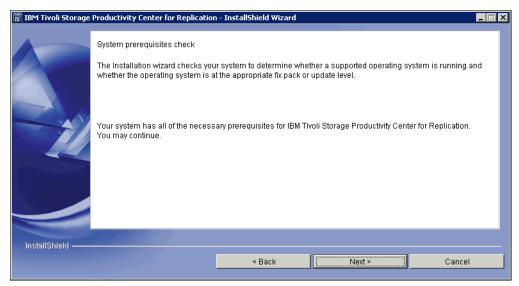


Figure 2-57 System prerequisites check running

If the system successfully passes the prerequisites check, the panel shown in Figure 2-58 opens. Click Next.

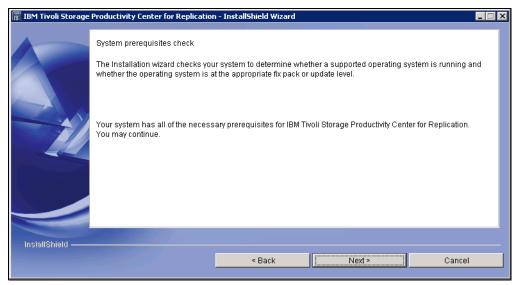


Figure 2-58 System prerequisites check passed

3. Review the license agreement. If you agree with the terms, select I accept the terms of the license agreement, and click Next as shown in Figure 2-59.

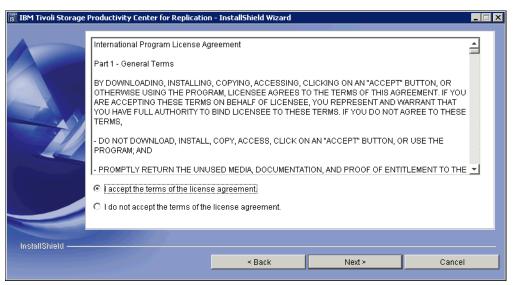


Figure 2-59 License agreement

4. On the next panel (Figure 2-60), you can select the directory where Tivoli Storage Productivity Center for Replication will be installed. A default location is displayed. You can accept it or change it based on your requirements. We install Tivoli Storage Productivity Center for Replication to the E: drive. Click **Next** to continue.

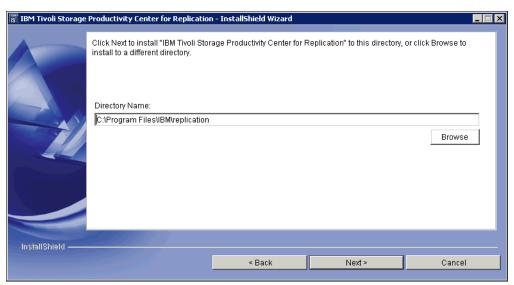


Figure 2-60 Destination Directory panel

5. In the next panel (Figure 2-61), select the Tivoli Storage Productivity Center for Replication user ID and Password. This ID is usually the system administrator user ID. If you are using local OS authentication and you want to enable the single sign-on feature for this user ID, you must provide the same credentials that you provided for the WebSphere Application Server Administrator (see step 9 on page 51).



Figure 2-61 User ID and password

Another user name: If you want to use another user ID, create it before beginning the installation and ensure that it has administrator rights.

6. The Default ports panel opens (Figure 2-62.). Ensure that the selected ports are available on the server and then click **Next**.

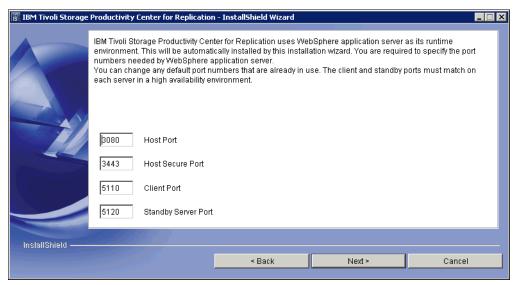


Figure 2-62 Tivoli Storage Productivity Center for Replication Ports panel

7. Review the settings, as shown in Figure 2-63, and click **Install** to start the installation.

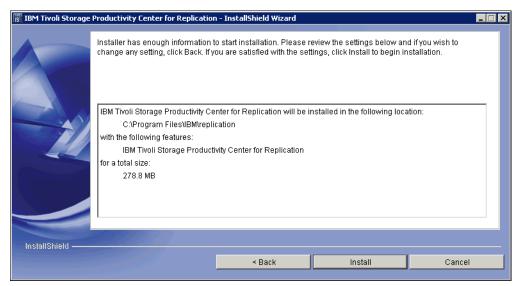


Figure 2-63 Tivoli Storage Productivity Center for Replication Settings panel

The installation of Tivoli Storage Productivity Center for Replication starts. Several messages about the installation process are shown, such as the one in Figure 2-64.

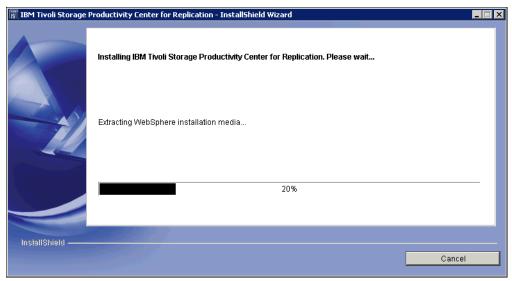


Figure 2-64 Tivoli Storage Productivity Center for Replication installation running

8. After the completion of the Tivoli Storage Productivity Center for Replication installation, a summary panel opens (Figure 2-65), which also indicates the web address to use to access the Tivoli Storage Productivity Center for Replication web-user interface. Click **Finish** to close the panel; the installation process resumes and the Tivoli Storage Productivity Center installation panel opens.

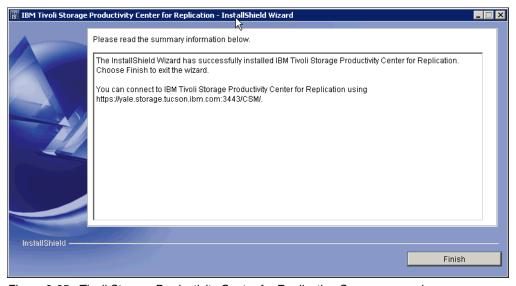


Figure 2-65 Tivoli Storage Productivity Center for Replication Summary panel

Licenses: Tivoli Storage Productivity Center for Replication is installed with no license. You must install the Two Site or Three Site Business Continuity (BC) license.

The Tivoli Storage Productivity Center for Replication is finished. You can now complete the installation of Tivoli Storage Productivity Center.

Completing the installation of the Tivoli Storage Productivity Center

After creation of the Tivoli Storage Productivity Center uninstaller, the summary information panel opens (Figure 2-66). Read and verify the information and click **Finish** to complete the installation.

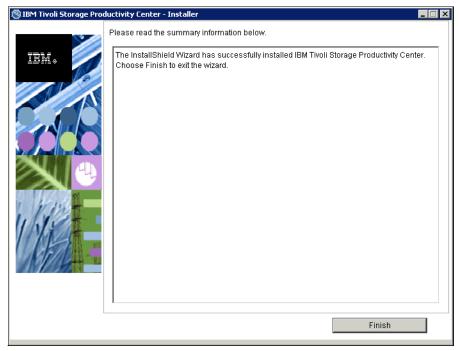


Figure 2-66 Component installation completion panel

Verifying the installation

At the end of the installation, we can use the Windows Services panel to verify that the Tivoli Storage Productivity Center services (Figure 2-67) have been installed.

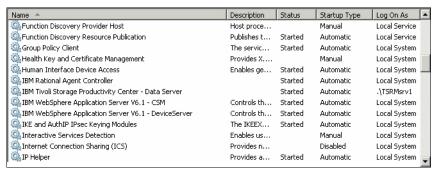


Figure 2-67 Windows service

The following services are related to Tivoli Storage Productivity Center:

- IBM Tivoli Storage Productivity Center Data Server
- IBM WebSphere Application Server v6.1 Device Server
- ▶ IBM WebSphere Application Server v6.1 CSM is the service related to Tivoli Storage Productivity Center for Replication

Another process, Tivoli Integrated Portal, is also present in the list of services (Figure 2-68).

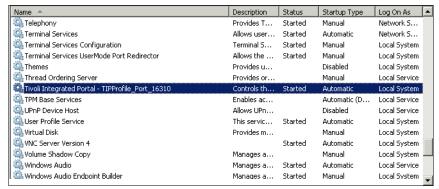


Figure 2-68 Tivoli Integrated Portal Process

All of the following services have to be present and started:

- ► Log files for Data server
- Log files for Device server
- Log files for GUI
- ▶ Log files for CLI

Log files for Data server

Check the logs for any errors or Java exceptions. On Windows, the default installation location is as follows:

c:\Program Files\IBM\TPC

The log files for the Data server are as follows:

- <InstallLocation>\TPC.log
- <InstallLocation>\log\data\install
- <InstallLocation>\log\install
- ► <InstallLocation>\data\log

Log files for Device server

Check the log files for any errors. The log files for the Device server are as follows:

- ► <InstallLocation>\TPC.log
- ► <InstallLocation>\log\device\install
- <InstallLocation>\device\log

Log files for GUI

Check the log files for any errors. The log files for the GUI are as follows:

- <InstallLocation>\TPC.log
- <InstallLocation>\log\gui\install
- ► <InstallLocation>\gui\log

Log files for CLI

Check the log files for any errors. The log files for the CLI are as follows:

- <InstallLocation>\TPC.log
- ► <InstallLocation>\log\cli\install

2.4.3 Agent installation

In this section, we present how to locally install Tivoli Storage Productivity Center agents.

Data agent or Fabric agent install

The Tivoli Storage Productivity Center V4.2 installation program does not support installation of the Data agent or Fabric agent. If you want to install the legacy Data agent or Fabric agent, you must have a previous Tivoli Storage Productivity Center installation program that supports installing the Data agent or Fabric agent.

Storage Resource agent installation

You typically install the Storage Resource agent by using the Tivoli Storage Productivity Center GUI. However, another possibility is to install it locally on a server through a command line. See 8.3, "Storage Resource Agent installation methods" on page 267 for more information and examples.

Depending on the decision of running the agent as a daemon or non-daemon service (on-demand service) and on the communication protocol that must be used, other parameters might be required.

The images of the Storage Resource agent are located on both Tivoli Storage Productivity Center images disks in the following location:

<DiskImage>\data\sra\windows

We navigate to the following directory:

<DiskImage>\data\sra\windows\bin

In our environment, the communication is between two Windows systems, so the default communication protocol used is Windows (SMB). We also want to run the agent as a non-daemon service. As a result, the command that we issue requires a minimum set of parameters, similar to the following set:

Agent -install -serverPort <serverport> -serverIP <serverIP> -installLoc <installLocation> -userID <userID> -password <password>

The meanings of these parameters and the values that we specified and use throughout this book are listed in Table 2-1.

Table 2-1	Storage Resource	agent instal	l parameters
-----------	------------------	--------------	--------------

Parameter	Explanation	Value
serverPort	The port of the Tivoli Storage Productivity Center Data server. The default value is 9549.	9549
serverIP	IP address or fully qualified DNS name of the server.	colorado.itso.ibm.com
installLoc	Location where the agent will be installed ^a .	c:\tpcsra
userID	The user ID defined on the agent system. This is the user ID t.hat the server can use to connect to the agent system	Administrator
password	Password for the specified User ID.	itso13sj

a. Be sure that when you specify a directory of where to install the Storage Resource agent, you do not specify an ending backslash (\) character. For example, do not specify C:\agent1\ because it causes the installation to fail.

Figure 2-69 shows a successful installation of the Storage Resource agent.

```
C:\WINDOW5\system32\cmd.exe

C:\TPC41$EDisk2\data\sra\windows\bin\Agent.exe -install -serverPort 9549 -server
IP colorado.itso.ibm.com -installLoc c:\tpcsra -userId Administrator -password i
tso13sj
AGT0383I Install completed successfully.

C:\TPC41$EDisk2\data\sra\windows\bin\_
```

Figure 2-69 Successful Storage Resource agent installation

To verify that the installation completed correctly from the Tivoli Storage Productivity Center GUI, log on to the Tivoli Storage Productivity Center GUI and go to **Administrative**Services

Data Sources

Data/Storage Resource Agents. The installed agent is now present in the list, as shown in Figure 2-70.

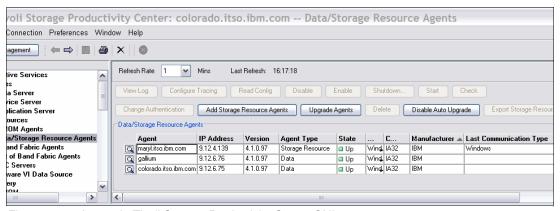


Figure 2-70 Agents in Tivoli Storage Productivity Center GUI

Tip: For the agent installed on server maryl.itso.ibm.com, the Agent Type column is Storage Resource and the Last Communication Type is Windows.

2.4.4 Disabling Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication

If you have installed Tivoli Storage Productivity Center V4.2 on a system with more than 4 GB of RAM but less than 8 GB we strongly suggest that you run only Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on that system. In this case, you must disable one of the two products. Also, if you have a powerful server but you plan to use only one of the two products, you can disable the other with the procedures we document in this section.

Disabling Tivoli Storage Productivity Center for Replication

To disable the Tivoli Storage Productivity Center for Replication server, complete these steps:

- Click to Start → Settings → Control Panel → Administrative Tools → Services. Right-click the IBM WebSphere Application Server V6.1 - CSM service.
- 2. Select **Properties**, as shown in Figure 2-71.

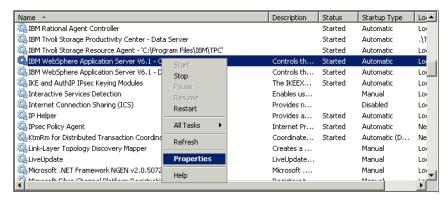


Figure 2-71 Tivoli Storage Productivity Center for Replication service

 On the panel shown in Figure 2-72, select **Disabled** under the Startup type menu, and click **Stop** in the Service Status section. When the service has been stopped, click **OK** to close this panel.

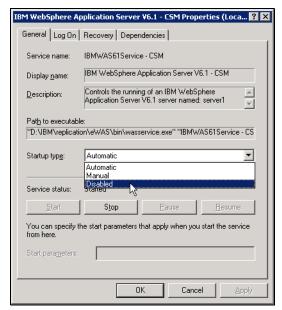


Figure 2-72 Service properties panel

Disabling Tivoli Storage Productivity Center

To disable the Tivoli Storage Productivity Center, complete the following steps:

- Click to Start → Settings → Control Panel → Administrative Tools → Services. Right-click the following service:
 - IBM WebSphere Application Server V6.1 DeviceServer
- 2. Select **Properties**, as shown in Figure 2-73.

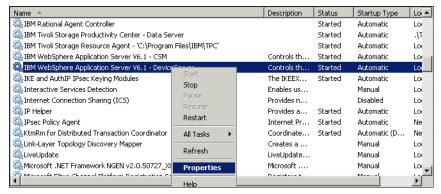


Figure 2-73 Services panel

 On the panel shown in Figure 2-74, select **Disabled** under the Startup type menu, and click **Stop** in the Service Status section. When the service stops, click **OK** to close this panel.

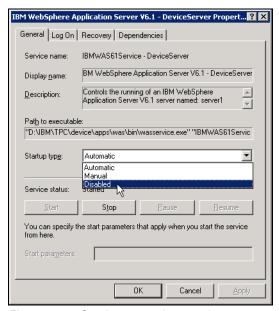


Figure 2-74 Service properties panel

- 4. Repeat the same procedure for the following services:
 - IBM Tivoli Storage Productivity Center Data server
 - IBM Tivoli Storage Resource agent < directory> if a Storage Resource agent is installed. Note that < directory> is where the Storage Resource agent is installed. The default is the <TPC_install_directory>\agent location.
- 5. Optional: You can also disable the following two services:
 - Tivoli Integrated Portal TIPProfile_Port_<xxxxx>. The <xxxxx> indicates the port that
 is specified during installation. The default port is 16310.
 - IBM ADE Service (Tivoli Integrated Portal registry)

Attention: Stop Tivoli Integrated Portal and IBM ADE Service only if no other applications are using these services and you are not using LDAP.

2.5 Applying a Tivoli Storage Productivity Center Upgrade

In this section, we cover the steps needed to apply a new build to Tivoli Storage Productivity Center on a Windows platform. Several steps remain the same for other platforms:

 Checking the version installed. To check your currently installed Tivoli Storage Productivity Center version, navigate to Help → About (Figure 2-75). A window opens (Figure 2-76), that shows the version, in this case the installed version is 4.2.2.62.



Figure 2-75 Tivoli Storage Productivity Center Help menu

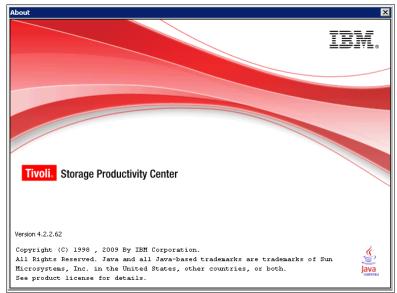


Figure 2-76 Tivoli Storage Productivity Center version installed

- 2. Prepare your environment; to upgrade, be sure that all GUIs are closed. There is no need to stop any Tivoli Storage Productivity Center service.
- 3. Run the installer. Double-click the **setup.exe** file to start the installation program. Select the language (Figure 2-77), and click **OK** to continue.



Figure 2-77 Installer, language selection

4. Review the license agreement (Figure 2-78) and, if you agree, accept the license and click **Next** to continue.

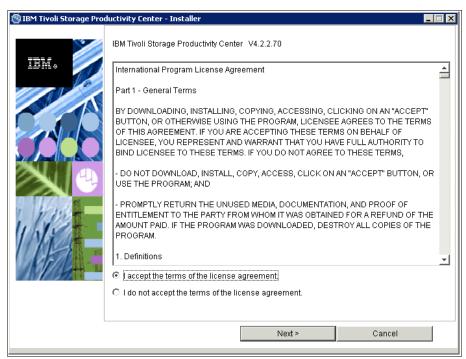


Figure 2-78 License agreement

5. Select Custom installation as the installation type (Figure 2-79). Click Next to continue.



Figure 2-79 Installation type

6. On the component selection window (Figure 2-80) all options are disabled so you cannot change them. Click **Next** to continue.



Figure 2-80 Component selection

7. Provide Database administrator information (Figure 2-81). The fields are filled in automatically. Click **Next** to continue.

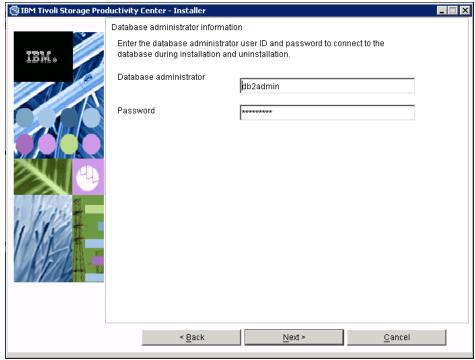


Figure 2-81 Database administrator information

8. Provide database schema information (Figure 2-82). The fields are automatically complete. Click **Next** to continue.

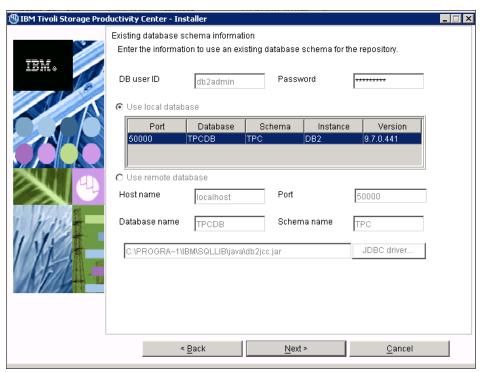


Figure 2-82 Database schema information

9. Confirm Data server, Device server, and Storage resource agent information (Figure 2-83).

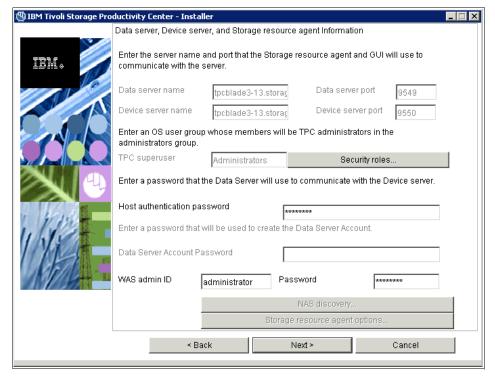


Figure 2-83 Confirm all information

10. Review the summary information (Figure 2-84) and click **Install** to proceed with the upgrade. During the installation process, the progress status is displayed.

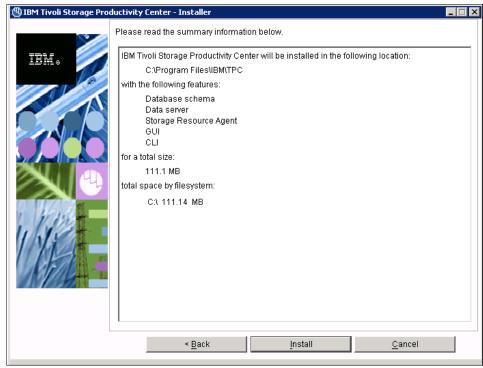


Figure 2-84 Summary

11. The installation wizard might not always shut down the device server service, and the installation can fail and issue the following message (Figure 2-85):

Cannot upgrade component Device Server



Figure 2-85 Error during installation

To avoid this error, kill the process that is using the following file:

C:\Program Files\IBM\TPC\device\apps\was

For that purpose, we used the "process explorer" utility as we show next:

a. You can download Process Explorer from the Microsoft website:

http://technet.microsoft.com/en-us/sysinternals/bb896653.aspx

It is an executable file (no installation required) that provides much more information from Windows processes than the already-included Windows Task Manager.

In the remaining steps, we show how it works by focusing on the issue we had with the device server process not being restarted.

b. After you download the utility, double-click its icon. A window opens, as shown in Figure 2-86. If your operating system is working on 64 bits, a new executable file named procexp64.exe is generated on the same path where the original resides.

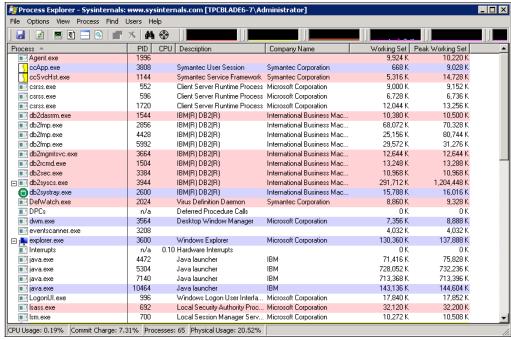


Figure 2-86 Process Explorer main window

c. To determine what process is causing the installation to fail, we must determine the image path for the process running. Click View → Select Columns (Figure 2-87).

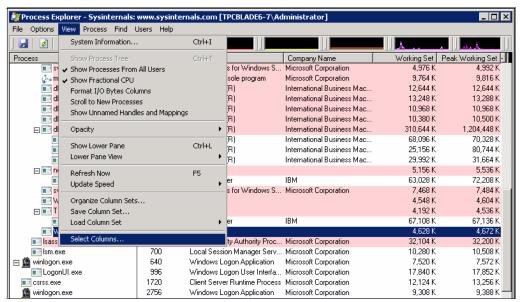


Figure 2-87 Select Columns from the View menu

d. On the Process Image tab, select the Image Path check box (Figure 2-88) to add a new column to Process Explorer main window where you are able to see the full image path of the each process running. Click OK to finish.

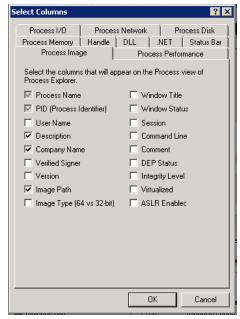


Figure 2-88 Select Columns window

e. Back in the Process Explorer window, scroll right to look at the image paths (Figure 2-89) and look for an entry like the one highlighted in the red box.

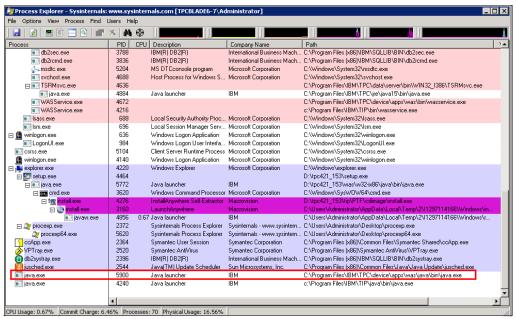


Figure 2-89 Path column

- f. Finally, the process highlighted is that one running on the path shown on the Tivoli Storage Productivity Center installation wizard. Kill the process either by pressing the Delete key or by right-clicking the process and selecting Kill Process. You should be able to continue with the installation by clicking Next on the Tivoli Storage Productivity Center installation wizard.
- 12. After the installation completes click **Finish** to exit the installation wizard (Figure 2-90).

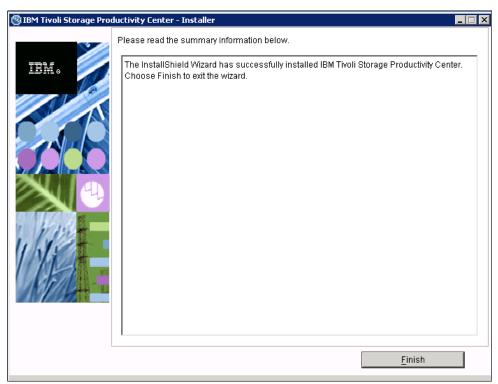


Figure 2-90 Installation completed window



Tivoli Storage Productivity Center installation on Linux

In this chapter, we show the installation of Tivoli Storage Productivity Center V4.2 on the Red Hat Linux platform. Of the available installation paths, Typical and Custom, we describe the Custom installation in our environment.

This chapter describes how to install IBM Tivoli Storage Productivity Center Standard Edition V4.2 and IBM Tivoli Storage Productivity Center for Replication V4.2 on 64-bit Red Hat Enterprise Linux 5 using the graphical interface. The prerequisite components (DB2 and Agent Manager) are installed prior to invoking the installation program.

This chapter also provides information about the preparation work required before installing the Tivoli Storage Productivity Center family.

3.1 Tivoli Storage Productivity Center installation on Linux

This section provides an overview of the installation, the media layout, and components.

3.1.1 Installation overview

To successfully install Tivoli Storage Productivity Center V4.2, use the following steps:

- 1. Check that the system meets the prerequisites. See 3.2, "Preinstallation steps for Linux" on page 81.
- 2. Install and configure all required prerequisite components.

 See "Prerequisite component for Tivoli Storage Productivity Center V4.2" on page 82.
- 3. Install Tivoli Storage Productivity Center database schema. See 3.3.2, "Creating the database schema" on page 101.
- Install Tivoli Storage Productivity Center server components.
 See 3.3.3, "Installing Tivoli Storage Productivity Center servers, GUI, and CLI" on page 107.

You can install Tivoli Storage Productivity Center family components using either Typical installation or Custom installation.

Typical installation

With the Typical installation, you install all the components of the Tivoli Storage Productivity Center on the local server in one step although you can still decide which components to install:

- Server: Data server, Device server, Replication Manager, and Tivoli Integrated Portal
- ► Clients: Tivoli Storage Productivity Center GUI
- Storage Resource agent

The "Typical" installation path is ideal for small and medium sized environments where you do not need to customize installation details and the defaults are sufficient.

The "Custom" path is preferred for large environments, or when you need to have greater control over installation choices, such as these:

- ► Size and location of the Tivoli Storage Productivity Center database, tables, logs, and so on.
- Greater flexibility with user accounts and passwords for components and services

Custom installation

With the Custom installation, you install parts of Tivoli Storage Productivity Center separately; The Custom installation method provides options so you can change default settings, such as user IDs, and directories. This is the installation method that we prefer.

Installing Tivoli Storage Productivity Center has the following installable components:

- ▶ Database Schema
- Data server and Device server
- Graphical user interface (GUI)
- ► Command-line interface (CLI)
- Storage Resource agent

Attention: Tivoli Storage Productivity Center for Replication is no longer a stand-alone application. Tivoli Storage Productivity Center Version 4.2 now installs Tivoli Integrated Portal and Tivoli Storage Productivity Center for Replication Version 4.2 during the server components installation process.

After Tivoli Storage Productivity Center Standard Edition is installed, the installation program will start the Tivoli Storage Productivity Center for Replication installation wizard.

The approximate time to install Tivoli Storage Productivity Center, including Tivoli Integrated Portal, is about 60 minutes. The approximate time to install Tivoli Storage Productivity Center for Replication is about 20 minutes.

3.1.2 Product code media layout and components

The web media consists of a disk image and a Storage Resource agent package:

- Disk 1 contains all Tivoli Productivity Center components:
 - Database Schema
 - Data server
 - Device server
 - GUI
 - CLI
 - Storage Resource agent
 - Tivoli Integrated Portal
 - Tivoli Storage Productivity Center for Replication

Tip: Disk 1 has 4 parts to it. All parts must be downloaded and extracted into the same directory.

Storage Resource agent package

The Storage Resource agent package contains the Tivoli Storage Productivity Center Storage Resource agent. It does not include a GUI installer.

To understand how this installation method works, see Chapter 8, "Storage Resource Agent" on page 263.

Physical media

The physical media included with the Tivoli Storage Productivity Center V4.2 product consists of a DVD and a CD. The DVD contains the Disk 1 parts described in "Product code media layout and components" on page 81. The physical media CD contains the Storage Resource agent package.

3.2 Preinstallation steps for Linux

Before installing Tivoli Storage Productivity Centeron Linux, analyze your environment to ensure that the system requirements have been met and that you have all the prerequisite components installed and configured.

3.2.1 Verifying system hardware and software prerequisites

For a detailed description of the system hardware and software prerequisites, and the latest platform support information, see the following website:

http://www.ibm.com/support/docview.wss?uid=swg27019380

3.2.2 Prerequisite component for Tivoli Storage Productivity Center V4.2

For Tivoli Storage Productivity Center V4.2, DB2 UDB Enterprise Server Edition is the only prerequisite component. The supported levels of IBM DB2 UDB Enterprise Server Edition are as follows:

- ► v9.1 (Fix Pack 2 or later)
- ► v9.5 (Fix Pack 6 not supported)
- v9.7 (Fix Pack 1 to 3a not supported

For the most current Tivoli Storage Productivity Center 4.2 platform support, see the following website:

http://www.ibm.com/support/docview.wss?uid=swg27019380

In this section, we show how to install the Tivoli Storage Productivity Center prerequisites on Linux. We perform a typical installation of DB2 v9.7 64-bit on Red Hat Enterprise Linux 5.

Be sure to verify that your system meets all minimum system requirements for installing the prerequisites, including adequate free disk space. See 3.2.1, "Verifying system hardware and software prerequisites".

Before beginning the installation, be sure to log on to your system as a local system user with root authority.

Migration: In this section, we are dealing with a clean installation of Tivoli Storage Productivity Center. If you are required to migrate your current Tivoli Storage Productivity Center environment to Version 4.2 from a previous release of Tivoli Storage Productivity Center, see *IBM Tivoli Storage Productivity Center Installation and Configuration Guide*, SC27-2337, Chapter 4: "Upgrading and migrating the IBM Tivoli Storage Productivity Center family."

http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/topic/com.ibm.tpc_V42.do c/fqz0_installguide_v42.pdf

3.2.3 DB2 installation: Using the GUI installation program

This topic describes how to install DB2 v9.5 Fix Pack 3a 64-bit DB2 on Linux using the GUI installation program.

Prerequisite: Before using the GUI to install DB2, you must have the X11 graphical capability installed. See Appendix C, "Configuring X11 forwarding" on page 815.

To install DB2, log on as a user with root authority, and then use the following procedures.

Accessing the installation media using the CD

Complete the following steps:

1. Create a mount point or choose an existing mount point. To create a mount point named /cdrom, enter the following command:

```
mkdir /cdrom
```

2. Insert the DB2 CD into the CD-ROM drive. Mount the CD-ROM file system at the desired mount point and run the following command:

```
mount -o ro /dev/cdrom /cdrom
```

3. Change to the directory where the CD-ROM is mounted:

cd /cdrom

Accessing the installation media using a downloaded image

Complete the following steps:

1. Create a temporary directory (for example, db2temp) to hold the DB2 installer .tar file and untarred (extracted) files. These files require from 2 - 3 GB of hard drive space.

```
mkdir /db2temp
```

- 2. Copy or download the DB2 installer into db2temp.
- 3. Change to the directory where you have stored the image, for example:

cd /db2temp

4. Extract the DB2 installer file, following the instructions supplied at the repository from which you downloaded the image, which might involve running the tar or gunzip commands, or a combination of both:

```
tar -xvzf v9.7_linuxx64_ese.tar.gz
```

5. Change to the installation directory, which you extracted from the image:

cd /db2temp/ese

Installing DB2

Complete the following steps:

- 1. Run the following command to verify that all necessary prerequisite packages are installed on the system:
 - ./db2prereqcheck

If during the prerequisite check you receive an error message, such as the one shown in Figure 3-1, you might need to install additional packages to satisfy DB2 dependencies before proceeding with the installation.

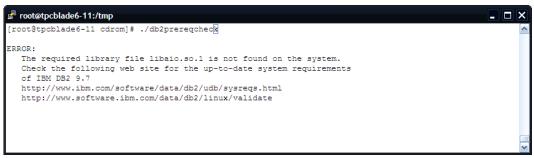


Figure 3-1 Error message indicating missing DB2 prerequisite packages

Go to the following address for more information about DB2 installation requirements for your specific platform:

http://www.ibm.com/software/data/db2/udb/sysreqs.html

- 2. Run the following command to execute the graphical installer:
 - ./db2setup

This command opens the DB2 Setup Launchpad (Figure 3-2).



Figure 3-2 DB2 Setup Launchpad

3. In the navigation tree, click **Install a Product** (Figure 3-3), and then select DB2 Enterprise Server Edition Version 9.7. Click **Install New** to proceed with the installation.

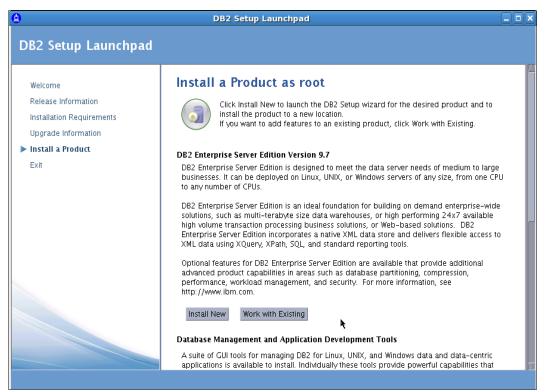


Figure 3-3 Click Install New to start the installation

Tip: When applying a DB2 fix pack, you must select the DB2 Enterprise Server Edition and not other options, as depicted in Figure 3-3.

4. The DB2 Setup wizard panel opens (Figure 3-4). Click **Next** to proceed.



Figure 3-4 DB2 Setup welcome message

 The next panel displays the software license agreement. Click Read non-IBM terms to display additional license information and, if you agree with all terms, click Accept and then Next to continue (Figure 3-5).

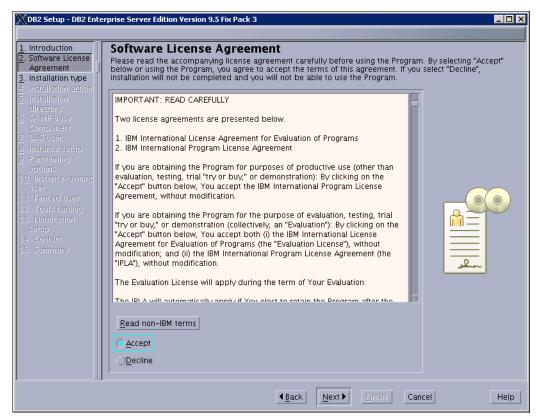


Figure 3-5 Software License Agreement

6. Accept the default installation type of **Typical** and click **Next** to continue (Figure 3-6).

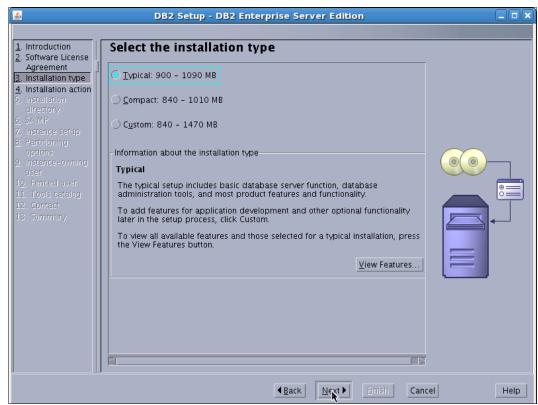


Figure 3-6 Select Typical installation type

7. In the next panel (Figure 3-7), accept the default Install DB2 Enterprise Server Edition on this computer and save my settings in a response file. Although not required, generate such a response file because it can greatly ease tasks such as documenting your work. Specify a valid path and the file name for the response file in the *Response file name* field. Click **Next** when you are ready.

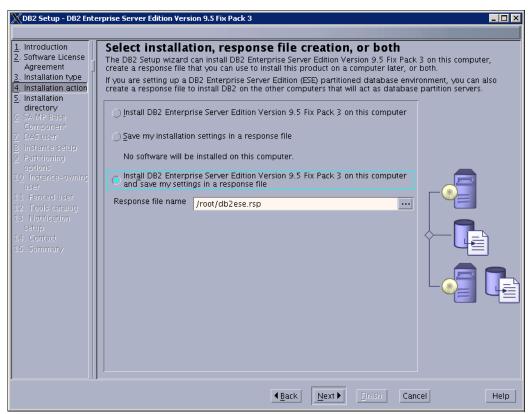


Figure 3-7 Select both installation and response file creation

8. The next panel (Figure 3-8) shows the default directory to be used as the installation folder. You can change the directory or accept the defaults. Make sure the installation folder has sufficient free space available, and then click **Next** to continue.



Figure 3-8 Select installation directory

9. If your system is an IBM System x® or System p®, you might see a panel titled "Install the IBM Tivoli System Automation for Multiplatforms Base Component (SA MP Base Component)." Because this component is not required by Tivoli Storage Productivity Center, so choose **Do not install SA MP Base Component** and click **Next**.

10. When you are prompted to set up a DB2 instance (Figure 3-9), accept the default to **Create a DB2 instance** and click **Next** to continue.

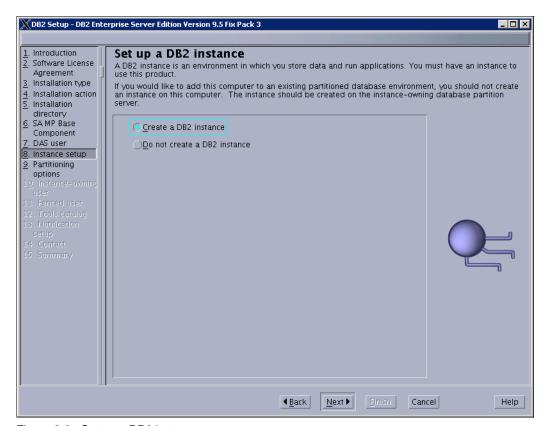


Figure 3-9 Set up a DB2 instance

11.In the next panel (Figure 3-10), accept the default to create a **Single partition instance** and click **Next** to continue.

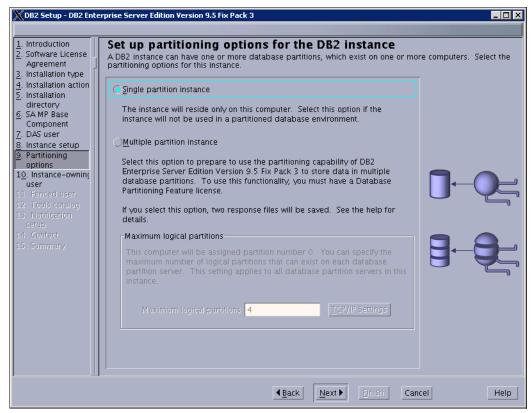


Figure 3-10 Choose to create a single partition instance

12. The next panel (Figure 3-11) asks you for user information for the DB2 instance owner. This user must have a minimal set of system privileges. Accept the default to create a **New user** and specify a password. Click **Next** when you are ready.

Space: The Tivoli Storage Productivity Center database repository will be stored in the home directory of the DB2 instance owner specified here. Make sure to place the user's home directory in a file system that has sufficient free space available; the /home directory is usually not large enough for database repositories.

In general, choose the file system that has the most available free space on your system to hold database repositories. If you are uncertain about the available file systems and their size, use the **df** -h command to get an overview.

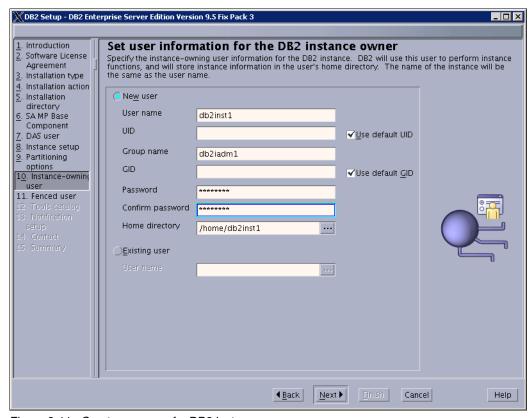


Figure 3-11 Create new user for DB2 instance owner

13. The last user you have to specify is the DB2 fenced user, which is used to execute user-defined functions (UDFs) and stored procedures. This user must have minimal system privileges also. Be sure you select **New user**, as shown in Figure 3-12. Specify a new password and click **Next** to continue.

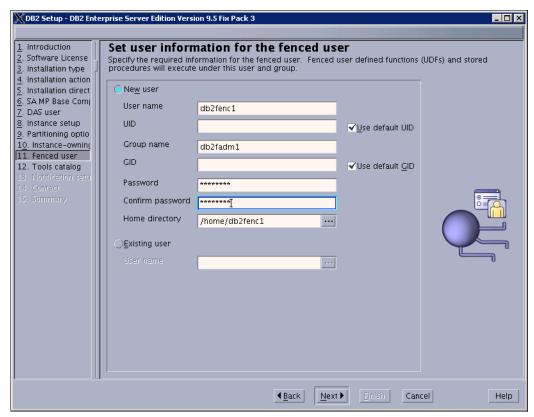


Figure 3-12 Create new user for DB2 instance owner

14. The next panel (Figure 3-13) prompts you to prepare the DB2 tools catalog. Because this component is not required by Tivoli Storage Productivity Center, click **Do not prepare the DB2 tools catalog**. Click **Next** to continue.

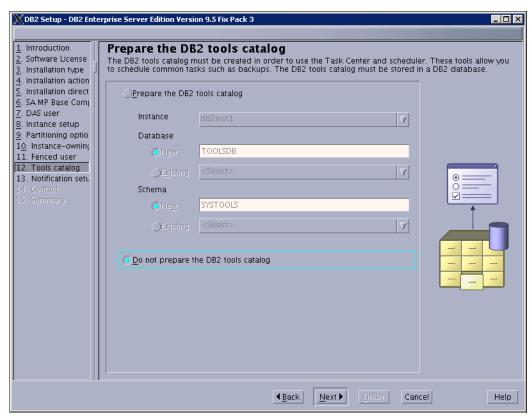


Figure 3-13 Choose not to prepare the DB2 tools catalog

15. In the next panel (Figure 3-14), you can specify a Notification SMTP (email) server. You can optionally specify an existing server or click **Do not set up your DB2 server to send notifications at this time**; a notification server can always be specified after the installation is finished. Make a choice and click **Next** to continue.

Tip: Configuring DB2 to send email notifications on errors and warning conditions can help resolve those conditions more quickly, thus improving overall stability and resiliency of the solution. This is an important factor in preventing unplanned outages.

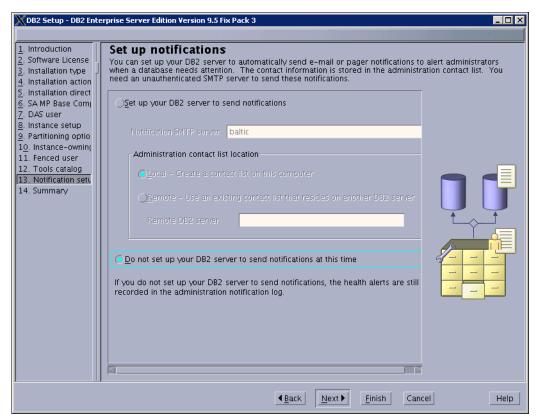


Figure 3-14 Optionally specify a notification server

16.A summary (Figure 3-15) lists what will be installed. Review all settings and, if you agree with them, click **Finish** to begin copying files.

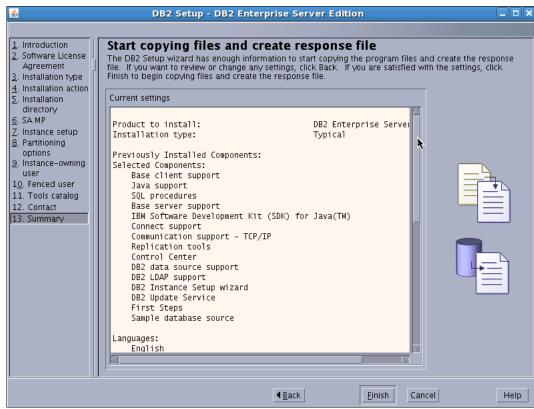


Figure 3-15 Installation summary

- 17.A progress panel is displayed as the installer copies the required files. Wait for the installation to complete.
- 18. When the installation successfully completes, the Setup Complete panel opens (Figure 3-16). Click **Finish** to close the panel.

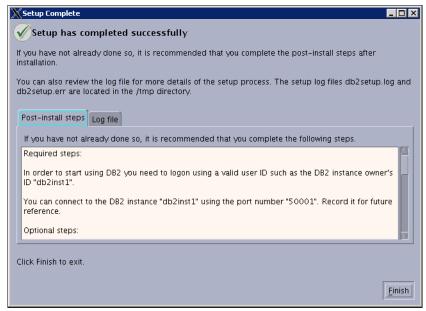


Figure 3-16 Setup Complete

19. After installation of DB2 completes, edit the /etc/group file by adding the root account to the db2iadm1 group. The db2iadm1 group line in the /etc/group file is as follows:

```
db2iadm1:x:102:root
```

20. After adding the root account to the db2iadm1 group, log out and log back in to allow the system to pick up this change.

3.2.4 Verifying that DB2 is installed correctly

The general steps to verify that DB2 has been installed properly are as follows:

- Create the SAMPLE database.
- Connect to the SAMPLE database.
- 3. Run a query against the SAMPLE database.
- 4. Drop the SAMPLE database.

Verify that DB2 has been installed properly:

1. Log on as a user with root authority.

Verification:

- ► After adding the root account to the db2iadm1 group, as outlined in "Installing DB2" on page 84, log out and log back in to allow the system to pick up this change.
- ▶ Before proceeding, check that root is a member of this group by issuing the id command. Make sure that the output line contains the db2iadm1 group; it looks similar to Figure 3-17.

```
root@tpcblade6-11:/tmp
[root@tpcblade6-11 tmp]# id
uid=0(root) gid=0(root) groups=0(root),1(bin),2(daemon),3(sys),4(adm),6(disk),10(wheel),105(db2iadm1)
[root@tpcblade6-11 tmp]# [
```

Figure 3-17 Verify that root is member of db2iadm1 group

- 2. To set the environment variables for the database instance, you need to source that the instance profile (db2profile) found in the instance user's home directory:
 - . /home/db2inst1/sqllib/db2profile

Attention: There is a space between the period (.) character and /home.

After setting the DB2 environment variables, verify the installed version of DB2 by issuing the db21eve1 command.

The output indicates which DB2 instance is currently being used, which code release is installed, and whether the selected DB2 instance is 32-bit or 64-bit, as shown in Figure 3-18.

Figure 3-18 Verify DB2 version and level

Important: Especially note whether the selected DB2 instance is 32-bit or 64-bit because this can greatly affect future installation steps.

4. Make sure that DB2 was started and is currently running by issuing the **db2start** command. If an error is issued, as shown in Figure 3-19, DB2 was already running when you issued the command. Otherwise DB2 is started now.

Figure 3-19 Verify that DB2 is running

5. Enter the **db2samp1** command to create the SAMPLE database. The results are similar to Figure 3-20.

Figure 3-20 Create sample database

Timing: This process can take several minutes to complete.

6. Enter the following commands to connect to the SAMPLE database, retrieve a list of all the employees that work in Department 20 (for example), and reset the database connection:

```
db2 connect to sample
db2 "select * from staff where dept = 20"
db2 connect reset
```

7. If all steps completed successfully, you can remove the SAMPLE database. Enter the following command to do so:

db2 drop database sample

The results look similar to Figure 3-21.

Figure 3-21 Verify DB2 installation

You have now successfully completed the DB2 installation.

3.3 Installing Tivoli Storage Productivity Center components

Now that the prerequisites have been installed, you can install the Tivoli Storage Productivity Center components.

3.3.1 Preparing for the installation

Before you begin the installation, be sure of the following requirements:

- Confirm that the correct version of DB2 is installed on your system.
- Be sure that the user IDs, which are required during the installation, are documented for reference.
- If you are considering the use of LDAP, ensure that you have all the correct information.
- See how to prepare the installation in IBM Tivoli Storage Productivity Center Installation and Configuration Guide, SC27-2337, Chapter 2, Installing the IBM Tivoli Storage Productivity Center family.

Tip: Install the Database Schema first. After that, install Data server, Device server, Tivoli Storage Productivity Center for Replication, and Tivoli Integrated Portal in a separate step.

If you install all the components in one step, and if any part of the installation fails for any reason (for example, space or passwords), the installation suspends and rolls back, uninstalling all the previously installed components.

Accessing the installation media using the CD

Complete the following steps:

1. Create a mount point or choose an existing mount point. To create a mount point called /cdrom, we enter the following command:

```
mkdir /cdrom
```

Insert the Tivoli Storage Productivity Center Disk 1 CD into the CD-ROM drive. Mount the CD-ROM file system at the desired mount point. Run the following command to achieve this:

```
mount -o ro /dev/cdrom /cdrom
```

3. Change to the installation directory where the CD-ROM is mounted, for example:

```
cd /cdrom
```

Accessing the installation media using a downloaded image

Complete the following steps:

1. Create a temporary directory (for example, tpctemp) to hold the Tivoli Storage Productivity Center installer tar files and untarred files. These files require 3 - 4 GB of hard drive space. Use the following command:

```
mkdir /tpctemp
```

- 2. Copy or download the Tivoli Storage Productivity Center installer into tpctemp.
- 3. Change to the directory where you have stored the image, for example:

```
cd /tpctemp
```

4. Extract the Tivoli Storage Productivity Center installer files, following the instructions supplied at the repository from which you downloaded the image, which might involve running the tar or gunzip commands, or a combination of both, for example:

```
tar -xvf TPC_4.2.1.108_SE_linux_ix86_disk1_part1.tar
tar -xvf TPC_4.2.1.108_linux_ix86_disk1_part2.tar
tar -xvf TPC_4.2.1.108_linux_ix86_disk1_part3.tar
tar -xvf TPC_4.2.1.108_linux_ix86_disk1_part4.tar
```

Tip: Be sure to extract all parts of Disk 1 into the same directory.

3.3.2 Creating the database schema

This topic provides information about how to create the database schema for use with Tivoli Storage Productivity Center.

Requirements:

- ▶ If you are using a remote database for Tivoli Storage Productivity Center, you must install the database schema on that computer first, after you have installed DB2. After this, you must install the Tivoli Storage Productivity Center server components on the other system, and choose to use a remote database connection.
- ► You must have the X11 graphical capability installed before using the GUI to install Tivoli Storage Productivity Center. See Appendix C, "Configuring X11 forwarding" on page 815

Installing the database schema

To install Tivoli Storage Productivity Center database schema, complete the following steps:

1. Log on as a user with root authority.

Verification:

- ► After adding the root account to the db2iadm1 group as outlined in "Installing DB2" on page 84, log out and log back in to allow the system to pick up this change.
- ▶ Before proceeding, check that root is a member of this group by issuing the id command. Make sure that the output line contains the db2iadm1 group.
- 2. To set the environment variables for the database instance, you must source the instance profile (db2profile) found in the instance user's home directory:
 - . /home/db2inst1/sqllib/db2profile

Attention: There is a space between . and /home.

- 3. Verify that DB2 was started and is currently running by issuing the **db2start** command. If an error is issued (Figure 3-19 on page 99), DB2 was already running when you issued the command. Otherwise DB2 is started now.
- 4. Change to the directory where you have extracted the Tivoli Storage Productivity Center Disk 1 software package, then launch the graphical installer by issuing the command:
 - ./setup.sh
- 5. Tivoli Storage Productivity Center installer is launched, prompting you to select an installation language (Figure 3-22). Choose a language and click **OK** to continue.



Figure 3-22 Select language

6. The International Program License Agreement is displayed. Read the license text and, if you agree with it, click I accept the terms of the license agreement (Figure 3-23). Click Next when you are ready to proceed with the installation.

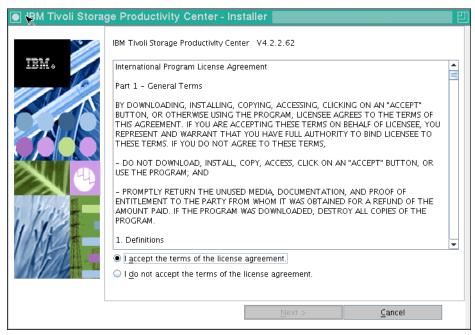


Figure 3-23 License agreement

7. In the next panel (Figure 3-24), select the **Custom installation** type. In addition, you can change the Tivoli Storage Productivity Center Installation Location to suite your requirements or accept the defaults. Make sure that the installation folder has sufficient free space available, and then click **Next** to continue.

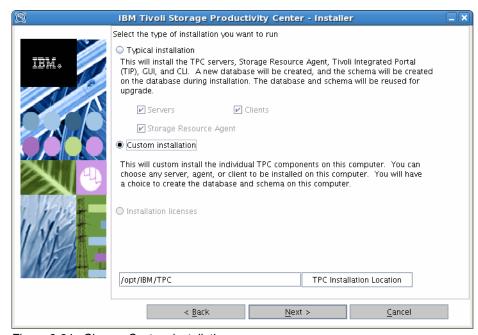


Figure 3-24 Choose Custom installation

8. In the next panel, (Figure 3-25), select one or more components to install. Remove all check marks except for **Create database schema** for now. Click **Next** to continue with the installation.

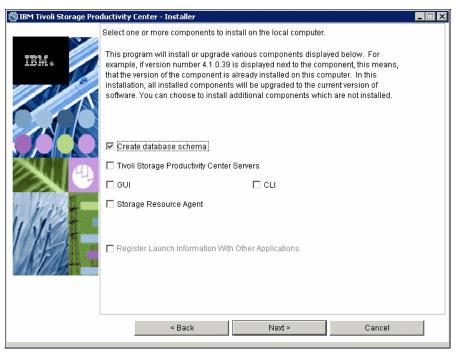


Figure 3-25 Select Create database schema component

9. The Database administrator information panel opens (Figure 3-26). Specify a user ID (with administrative database authority) in the **Database administrator** field, such as db2inst1, specify a password, and click **Next** to continue.

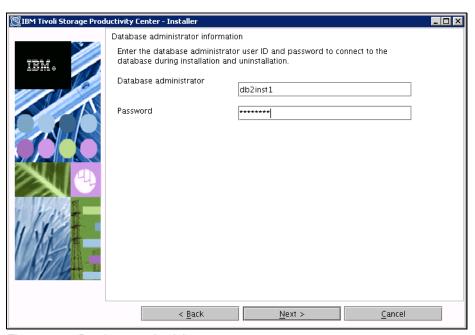


Figure 3-26 Database credentials

10. In the next panel (Figure 3-27), enter the administrative user ID (in our case, db2inst1) and password, again as *DB user ID*. Select **Create local database**.

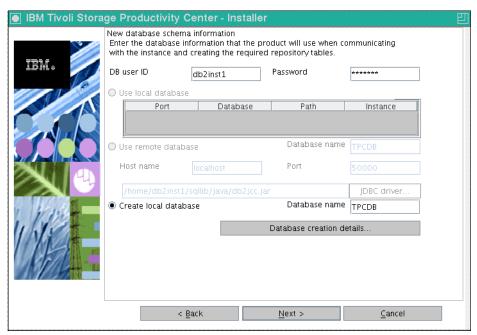


Figure 3-27 Choose Create local database

You can click **Database creation details** to verify additional details, as shown in Figure 3-28. Do not change the default values unless you are a knowledgeable DB2 administrator. Click **OK** to close the window, and then click **Next** to proceed with the installation.

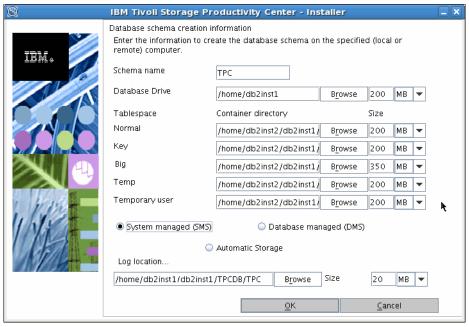


Figure 3-28 Database schema creation

Character length: The Tivoli Storage Productivity Center schema name cannot be longer than eight characters.

11. Figure 3-29 shows the summary information panel. Review the information that you have provided for the database schema installation. If you are in agreement that all data entered is correct, click **Install**.

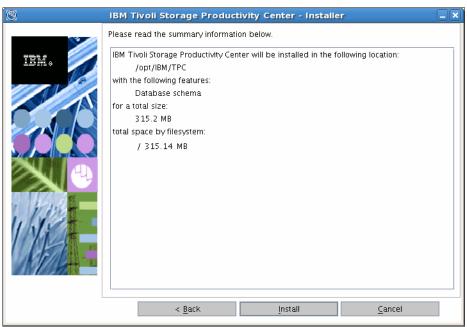


Figure 3-29 Summary information

12. The progress panel is displayed. Wait for the installation to finish; the results panel looks like Figure 3-30. Click **Finish** to exit the graphical installer.

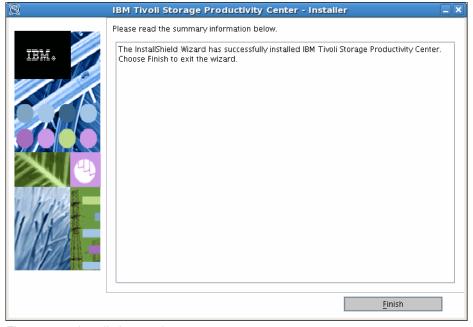


Figure 3-30 Installation results

Verifying the database schema installation

Verify the installation by checking that you have the database named *TPCDB*, as follows:

- 1. Source the DB2 profile:
 - . /home/db2inst1/sqllib/db2profile
- 2. Verify creation of the TPCDB database by issuing the following command:

```
db2 list db directory
```

The command lists all databases that exist, as shown in Figure 3-31.

```
[db2inst1@madden07 ~]$ db2 list db directory

System Database Directory

Number of entries in the directory = 1

Database 1 entry:

Database alias = TPCDB I

Database alias = TPCDB I

Database name = TPCDB = TPCDB I

Local database directory = /home/db2inst1

Database release level = d,00

Comment = Indirect

Catalog database partition number = 0

Alternate server hostname = Alternate server port number = 0
```

Figure 3-31 Verify database creation

3.3.3 Installing Tivoli Storage Productivity Center servers, GUI, and CLI

After you have completed creating the database schema, you are ready to install the following Tivoli Storage Productivity Center components:

- Data server
- Device server
- ► GUI
- ► CLI

In addition to these components, two additional components are installed by default:

- ► Tivoli Integrated Portal
- ► Tivoli Storage Productivity Center for Replication

Next we describe how to complete the installation process.

Prerequisite: X11 allows the use of a GUI. You must have the X11 graphical capability installed before using the GUI to install Tivoli Storage Productivity Center. See Appendix C, "Configuring X11 forwarding" on page 815.

Complete the installation process as follows:

- 1. Log on as a user with root authority.
- 2. To set the environment variables for the database instance, you must source the instance profile (db2profile) found in the instance user's home directory:
 - . /home/db2inst1/sqllib/db2profile

Attention: There is a space between . and /home.

3. Make sure that DB2 was started and is currently running by issuing the **db2start** command. If an error message is issued, as shown in Figure 3-19 on page 99, DB2 was already running when you issued the command. Otherwise the command is started now.

- 4. Change to the directory where you have extracted the Tivoli Storage Productivity Center Disk 1 software package, and then launch the graphical installer by issuing the following command:
 - ./setup.sh
- 5. Tivoli Storage Productivity Center installer is launched, prompting you to select an installation language (Figure 3-32). Choose a language and click **OK** to continue.



Figure 3-32 Select language

6. The international license agreement is displayed. Read the terms and, if you agree with it, click I accept the terms of the license agreement (Figure 3-33). Click Next.

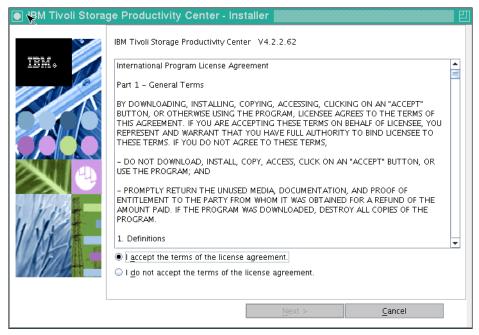


Figure 3-33 License agreement

7. In the next panel (Figure 3-34), select Custom installation, and then click Next.

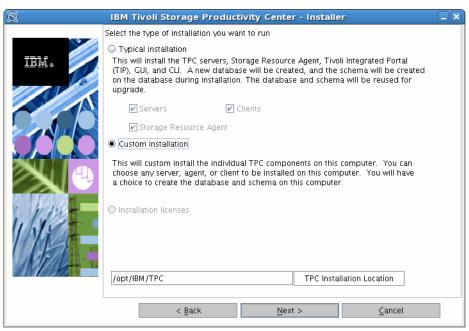


Figure 3-34 Choose Custom Installation

8. In the next panel (Figure 3-35), clear all the check marks except for **Tivoli Storage Productivity Center Servers**, **GUI**, and **CLI** components.

The **Create database schema** check box is disabled because it was installed previously. Click **Next** to continue with the installation.

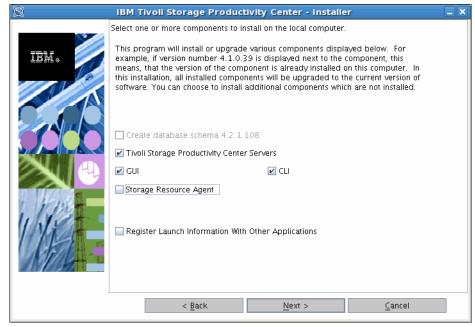


Figure 3-35 Select Servers, GUI, and CLI

Attention: Preferably, *do not* install the Storage Resource agent at this time. Installing any Storage Resource agent by the installer requires you to also uninstall the Storage Resource agent using the installer. Therefore, in most cases, later using the Tivoli Storage Productivity Center GUI to deploy agents (instead of using the installer) is the more flexible approach.

Because we do not plan to install the Storage Resource agent at this time, there is no need to Register with the agent manager; we perform this step subsequently.

9. If you are running the Tivoli Storage Productivity Center installation on a system with at least 4 GB but less than the preferable 8 GB of RAM, a warning message about memory size is issued (Figure 3-36). To ignore this message and continue with the installation, click **OK**.

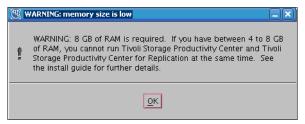


Figure 3-36 Memory size warning

RAM: 8 GB of RAM is the minimum memory requirement to run *both* Tivoli Storage Productivity Center and Tivoli and Tivoli Storage Productivity Center for Replication.

If you have less than 8 GB of RAM, you have to run *either* Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication because of system load. To do that, you must disable Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication after installation. See 2.4.4, "Disabling Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication" on page 67.

10.The Database administrator information panel opens (Figure 3-37). The database administrator user and password are automatically filled in. The reason is because we previously used it to create the database schema. Click **Next** to continue.

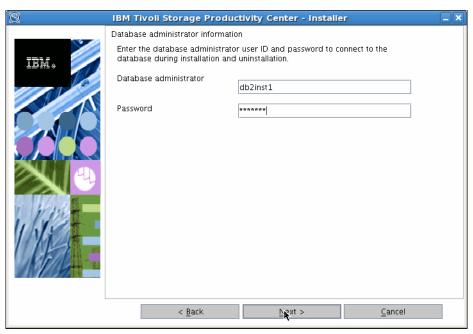


Figure 3-37 Database credentials

11. The Existing database schema information panel opens (Figure 3-38). Because we already installed the database schema, nothing can be changed here. Click **Next** to continue with the installation.

If you want to use a remote database on another system, you must install the Tivoli Storage Productivity Center schema component on this machine first, following the procedure documented in 3.3.2, "Creating the database schema" on page 101. After installing the Tivoli Storage Productivity Center server components, select the **Use remote database** option, and specify the host name of the server running the DB2 Manager.

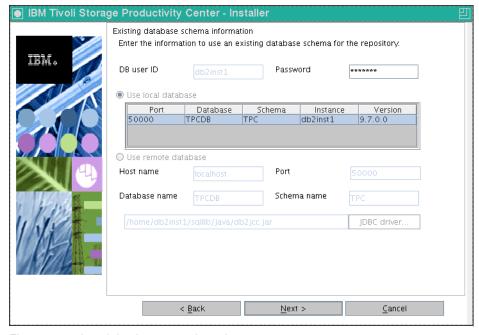


Figure 3-38 Local database preselected

12. If you selected to use a remote database, a message is issued (Figure 3-39) to ensure that the remote DB2 instance is running before you proceed. Click **OK**.

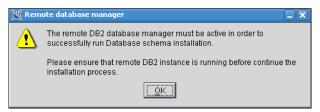


Figure 3-39 Ensure that DB2 is running on the remote system

- 13. The panel shown in Figure 3-40 requires the following input:
 - Data Server Name:

Enter the fully qualified host name of the Data server.

Data Server Port:

Enter the Data server port. The default is 9549.

– Device Server Name:

Enter the fully qualified host name of the Device server.

Device Server Port:

Enter the Device server port. The default is 9550.

TPC superuser:

Enter an operating system group name to associate with the Tivoli Storage Productivity Center superuser role. This group must exist in your operating system before you install Tivoli Storage Productivity Center. Membership in this group provides full access to the Tivoli Storage Productivity Center product. You can assign a user ID to this group on your operating system and log on to the Tivoli Storage Productivity Center GUI using this user ID.

If you click **Security roles**, the Advanced security roles mapping panel opens. You can assign an operating system group for each Tivoli Storage Productivity Center role with which you want to make an association, so you can have separate authority IDs to do various Tivoli Storage Productivity Center operations. The operating system group must exist before you can associate a Tivoli Storage Productivity Center role with it. Except for the superuser role, you do not have to assign security roles at installation time; you can assign these roles after you have installed Tivoli Storage Productivity Center.

Override: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the values you enter for LDAP Tivoli Storage Productivity Center Administrator groups override the values you entered here for the Tivoli Storage Productivity Center superuser.

You can record information that is used in the component installation (such as user IDs, passwords, and storage subsystems) in the worksheets in Appendix B, "Worksheets" on page 803.

Host Authentication Password:

This password is the password that is used for the Fabric agents to communicate with the Device server. This password must be specified when you install the Fabric agent.

Data Server Account Password:

This password is not required for Linux installations; it is only required for Windows.

 WebSphere Application Server (WebSphere Application Server) Admin ID and Password:

This information is the user ID and password that required by the Device server to communicate with the embedded WebSphere.

You can use any existing user ID here, such as the dasusr1 ID created upon DB2 installation. The WebSphere Application Server admin ID does not need to have any operating system privileges. The user ID will be used for the local Tivoli Integrated Portal (TIP) Administrator ID.

If you click the **NAS discovery** button, the NAS discovery information panel is displayed. You can enter the NAS filer login default user name and password and the SNMP communities to be used for NAS discovery. You do not have to assign the NAS discovery information at installation time, you can configure it after you have installed Tivoli Storage Productivity Center.

Important: Ensure that you record all passwords that are used during the installation of Tivoli Storage Productivity Center.

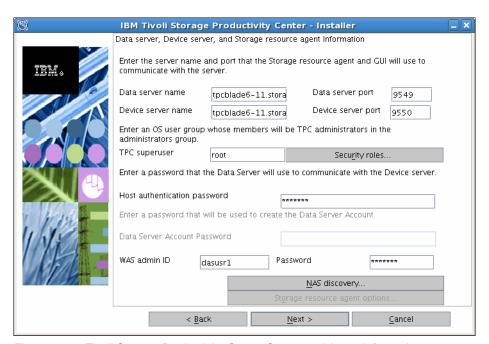


Figure 3-40 Tivoli Storage Productivity Center Server and Agent information

When you are ready, click **Next** to continue.

14. The Tivoli Integrated Portal panel opens (Figure 3-41). You can select to install a new version of Tivoli Integrated Portal or use an already existing installation on the local machine.

Tivoli Integrated Portal will use 10 port numbers starting from the one specified in the Port field (referred to as the *Base Port*):

- base port
- base port+1
- base port+2
- base port+3
- base port+5
- base port+6
- base port+8
- base port+10
- base port+12
- base port+13

The TIP Administrator ID and Password fields are prefilled with the WebSphere Application Server admin ID and password that are specified in the previous step (Device server installation).

Click Next to continue.

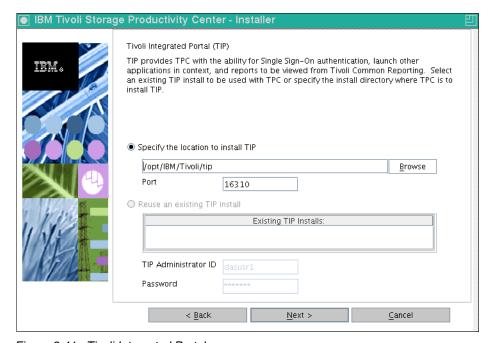


Figure 3-41 Tivoli Integrated Portal

Important: Tivoli Storage Productivity Center Version 4.2 supports a Tivoli Integrated Portal instance that is exclusively used only by Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication, but no other application exploiting Tivoli Integrated Portal.

15. The authentication selection panel opens (Figure 3-42). This panel refers to the authentication method that will be used by Tivoli Storage Productivity Center to authenticate the users.

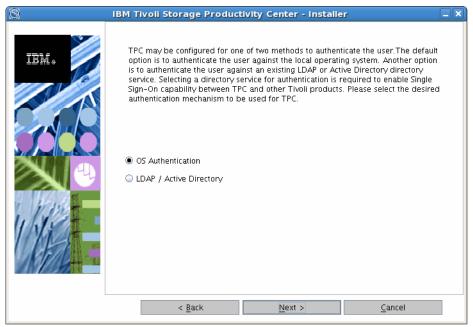


Figure 3-42 Authentication selection

If you already have a valid Tivoli Integrated Portal instance on the system and it uses either OS-based or LDAP-based authentication, Tivoli Storage Productivity Center will use that existing authentication method.

Otherwise, select the authentication method to use:

- OS Authentication
 - This method uses the operating system of the Tivoli Storage Productivity Center server for user authentication.
- LDAP/Active Directory
 - If you select LDAP or Microsoft Active Directory for authentication, you must have LDAP or Active Directory installed already.

If you select this method, additional panels are displayed to configure this authentication method. See 3.3.3, "Installing Tivoli Storage Productivity Center servers, GUI, and CLI" on page 107 for additional details.

Click **Next** to proceed.

16. The summary information panel opens (Figure 3-43). Review the information, and then click **Install** to continue.

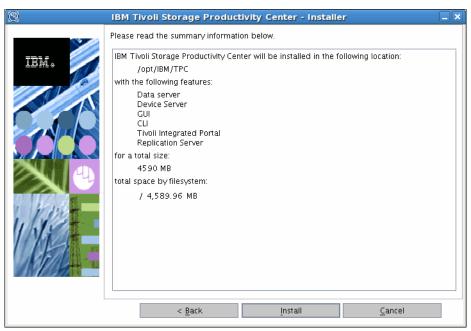


Figure 3-43 Summary information

- 17. The progress window is displayed as Tivoli Storage Productivity Center is installed. Wait for the installation to complete.
- 18. After the Tivoli Storage Productivity Center Data server, Device server, GUI, and CLI installation are complete, the Installing TIP message is displayed (Figure 3-44). Wait for the Tivoli Integrated Portal installation to finish.

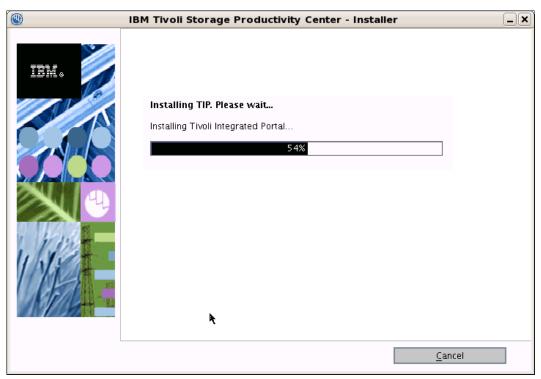


Figure 3-44 Tivoli Integrated Portal installation

After the Tivoli Integrated Portal installation has completed, the Tivoli Storage Productivity Center for Replication installation is launched in a separate window. The Tivoli Storage Productivity Center installation is temporarily suspended in the background and the Tivoli Storage Productivity Center for Replication panel is displayed as seen in Figure 3-45.

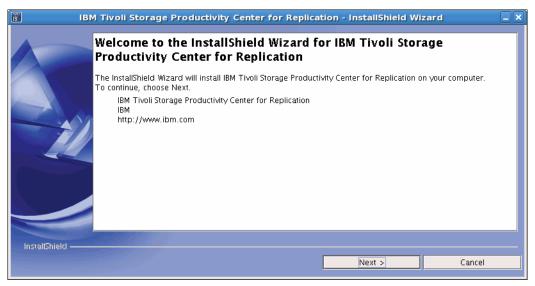


Figure 3-45 Tivoli Storage Productivity Center for Replication installation

Installing Tivoli Storage Productivity Center for Replication

Complete the following steps:

1. The Welcome panel opens (Figure 3-45). Click Next.

Important: If you are not planning to use Tivoli Storage Productivity Center for Replication and you attempt to cancel or bypass the installation, an interruption of the installation process is the result, which invokes a complete Tivoli Storage Productivity Center installation rollback.

2. The system prerequisites check panel opens (Figure 3-46). The wizard checks that the operating system meets all prerequisite requirements and fix packs are installed.

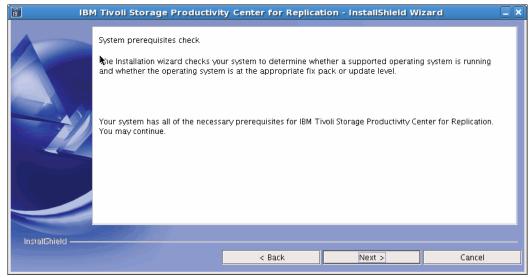


Figure 3-46 System check

If the system passes the check (Figure 3-47), click Next.

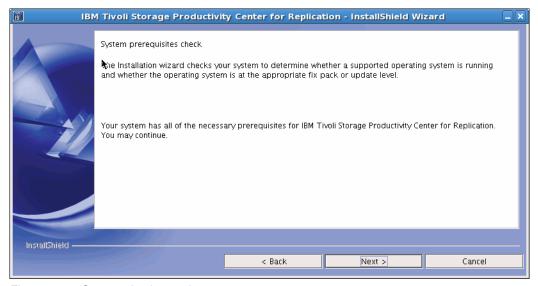


Figure 3-47 System check complete

3. Read the license agreement (Figure 3-48). If you agree with it, select I accept the terms of the license agreement, and click Next.

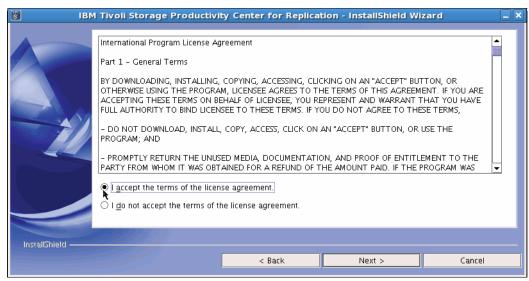


Figure 3-48 License Agreement

4. Specify the Directory Name where you want to install Tivoli Storage Productivity Center for Replication. You can choose a directory either by changing the location or by accepting the default directory. Figure 3-49 shows an example. Make sure that the installation folder has sufficient free space available, then click **Next** to continue with the installation.

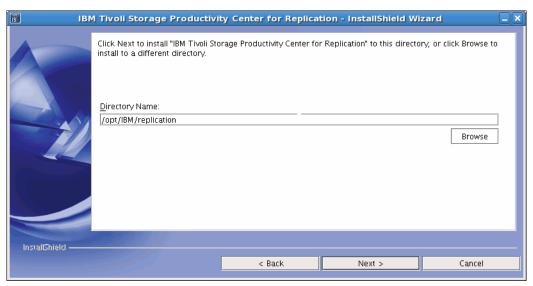


Figure 3-49 Directory Name

5. The Tivoli Storage Productivity Center for Replication Administrator user panel opens (Figure 3-50). Enter the user ID and password that will be used as Tivoli Storage Productivity Center for Replication administrator. This user must already exist in the operating system and have administrator rights, such as the root account. Click **Next** to continue.

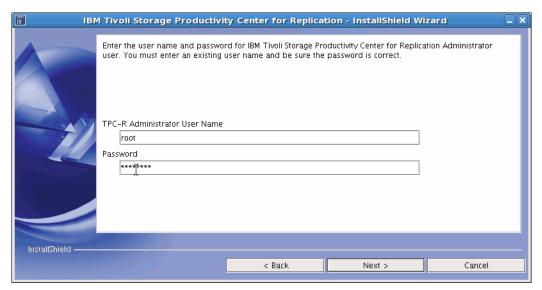


Figure 3-50 Tivoli Storage Productivity Center for Replication user ID and password

Tip: If you prefer to use another user, you must have already created it and ensure that it has administrator rights.

The default WebSphere Application Server ports panel opens (Figure 3-51). Accept the defaults and click **Next** to continue.

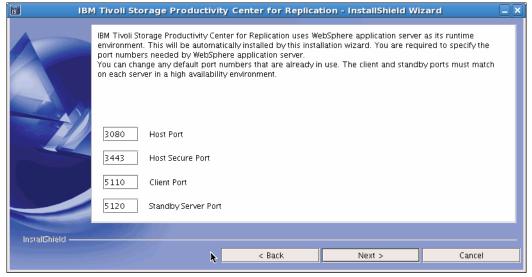


Figure 3-51 Default ports

7. Review the settings in the installation summary list (Figure 3-52). If necessary, make changes by clicking **Back**. Otherwise, click **Install**.

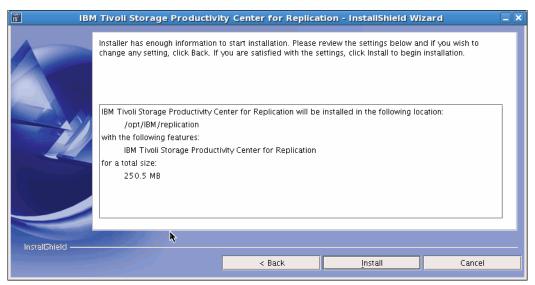


Figure 3-52 Tivoli Storage Productivity Center for Replication installation summary

8. The Tivoli Storage Productivity Center for Replication installation progress panel is displayed (Figure 3-53). Wait for the installation to finish.

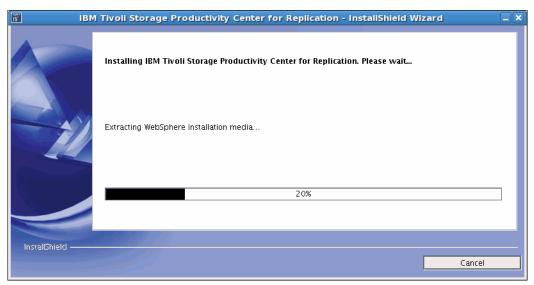


Figure 3-53 Tivoli Storage Productivity Center for Replication progress panel

9. The Tivoli Storage Productivity Center for Replication installation results are displayed (Figure 3-54). Notice the URL that is listed for connecting to Tivoli Storage Productivity Center for Replication. Click **Finish**.

Licenses: Tivoli Storage Productivity Center for Replication is installed with FlashCopy as the only licensed service. You must install the Two Site or Three Site Business Continuity (BC) license to use synchronous Metro Mirror and asynchronous Global Mirror capabilities.

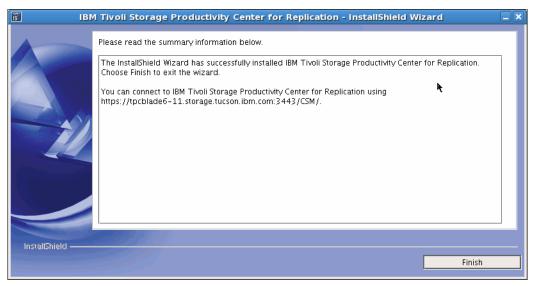


Figure 3-54 Tivoli Storage Productivity Center for Replication installation results

10. After the Tivoli Storage Productivity Center for Replication installation has completed, the Tivoli Storage Productivity Center installer continues creating the uninstaller (Figure 3-55). Wait for the installation to complete.

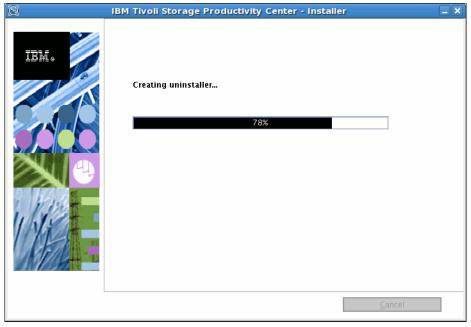


Figure 3-55 Creating uninstaller

11. After the installation has finished, a summary is displayed (Figure 3-56). Read and verify the information and click **Finish** to complete the installation.

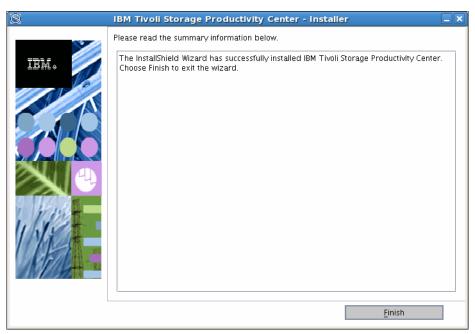


Figure 3-56 Summary Information

Verifying the Tivoli Storage Productivity Center server installation

At the end of the installation, a good approach is to make sure that all the components have been installed successfully and that Tivoli Storage Productivity Center is in good working order.

To test this situation on Linux, we launch the Tivoli Storage Productivity Center GUI. In Tivoli Storage Productivity Center, we then confirm that all services are started and running.

Complete the following steps:

- 1. If you installed Tivoli Storage Productivity Center to the default location, the following command launches the Tivoli Storage Productivity Center GUI on Linux:
 - /opt/IBM/TPC/gui/TPCD.sh
- 2. Log on to Tivoli Storage Productivity Center using a user ID that is mapped to the Tivoli Storage Productivity Center superuser role. If you used the defaults during installation, the *root* user is mapped to this role.
- 3. From the Navigation Tree, expand **Administrative Services** → **Service** → **Data Server** and **Device Server**. All nodes within these branches are marked green, as illustrated in Figure 3-57.

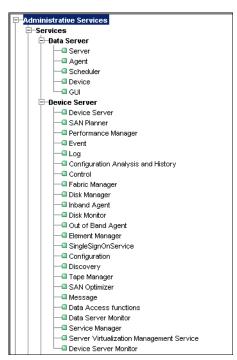


Figure 3-57 Data and Device server services

You have now successfully completed Tivoli Storage Productivity Center server installation.

Tivoli Storage Productivity Center installation on AIX

In this chapter, we show the installation steps of the Tivoli Storage Productivity Center V4.2 on the AIX platform.

4.1 Tivoli Storage Productivity Center installation on AIX

In this chapter, we describe how to install the Tivoli Storage Productivity Center Standard EditionV4.2 and Tivoli Storage Productivity Center for Replication V4.2. The prerequisite component (DB2) is installed prior to invoking the installation program.

Furthermore, in this section we provide information about the preparation work required before installing the Tivoli Storage Productivity Center family. This section provides an overview of the installation, the media layout, and components.

4.1.1 Installation overview

To get Tivoli Storage Productivity Center V4.2 to work, use the following steps:

- 1. Check that the system meets the prerequisites. See 4.2, "Preinstallation steps for AIX" on page 127.
- 2. Install and configure all required prerequisite components. See 4.3, "Installing the prerequisite for AIX" on page 128.
- Install Tivoli Storage Productivity Center components.
 See 4.4, "Installing Tivoli Storage Productivity Center components" on page 134.

You can install Tivoli Storage Productivity Center using either Typical installation or Custom installation.

With Custom installation, you see what components are being installed and where they are being installed, and you can customize your environment by installing separate components, versions, supply various passwords for user IDs, and change the default installation directories if required.

In our case, we install Tivoli Storage Productivity Center using the Custom installation option.

Attention: Starting with Tivoli Storage Productivity Center V4.1, Tivoli Storage Productivity Center for Replication is no longer a stand-alone application. Tivoli Storage Productivity Center V4.2 installs Tivoli Integrated Portal and Tivoli Storage Productivity Center for Replication V4.2.

The Custom installation of Tivoli Storage Productivity Center has the following installable components:

- ► Database schema
- ► Data server and Device server
- Graphical User Interface (GUI)
- Command Line Interface (CLI)

After Tivoli Storage Productivity Center is installed, the installation program starts the Tivoli Storage Productivity Center for Replication installation wizard.

4.1.2 Product code media layout and components

In this section, we describe the contents of the product media at the time of writing. The media content differs, depending on whether you are using the web images or the physical media included with the Tivoli Storage Productivity Center V4.2 package.

Passport Advantage and web media content

The web media consists of a disk image and a Storage Resource agent package:

- Disk 1 contains all Tivoli Productivity Center components:
 - Database Schema
 - Data server
 - Device server
 - GUI
 - CLI
 - Storage Resource agent
 - Tivoli Integrated Portal
 - Tivoli Storage Productivity Center for Replication

Tip: Disk 1 has 4 parts to it. All parts must be downloaded and extracted into the same directory.

The Storage Resource agent package contains Tivoli Storage Productivity Center Storage Resource agent. It does not include a GUI installer.

To understand how this installation method works, see Chapter 8, "Storage Resource Agent" on page 263.

Physical media

The physical media included with the Tivoli Storage Productivity Center V4.2 product consists of a DVD and a CD. The DVD contains the Disk 1 parts described in "Passport Advantage and web media content" on page 127. The physical media CD contains the Storage Resource agent package.

4.2 Preinstallation steps for AIX

Before installing Tivoli Storage Productivity Center V4.2 on AIX, you need to analyze your environment to ensure that the system requirements have been met and that you have all the prerequisite components installed and configured. Review this section for prerequisite information.

4.2.1 Verifying system hardware prerequisites

For the hardware prerequisites, see the following website:

http://www.ibm.com/support/docview.wss?uid=swg27019380

The Tivoli Storage Productivity Centerserver requires 8 GB of RAM. If you have at least 4 GB but less than 8 GB of RAM, you can still install Tivoli Storage Productivity Centerand Tivoli Storage Productivity Centerfor Replication. However, a warning message is issued during installation.

If you have less than 8 GB of RAM, you have to run only Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication because of system load. To do that, you must disable Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication after installation.

For installations on AIX, you need a total of 6 GB of free disk space.

- ▶ 2.25 GB for the /tmp directory
- ▶ 3 GB for the /opt directory
- ▶ 250 MB in the /home directory
- ▶ 10 KB of free space in /etc directory
- ▶ 200 MB in the /usr directory
- ► 50 MB in the /var directory

4.2.2 Verifying system software prerequisites

For the software prerequisites, see the following website:

http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=/com.ibm.tp
c V41.doc/fqz0 r sw requirements.html

4.2.3 Prerequisite component for Tivoli Storage Productivity Center V4.2

For Tivoli Storage Productivity Center V4.2, DB2 UDB Enterprise Server Edition is the only prerequisite component. The supported levels of IBM DB2 UDB Enterprise Server Edition are as follows:

- ▶ v9.1 (Fix Pack 2 or later)
- ▶ v9.5 (Fix Pack 6 not supported)
- ▶ v9.7 (Fix Pack 1 to 3a not supported)

For the most current Tivoli Storage Productivity Center 4.2 platform support, see this website:

http://www.ibm.com/support/docview.wss?uid=swg27019380

4.3 Installing the prerequisite for AIX

In this section, we show how to install the Tivoli Storage Productivity Center prerequisites on AIX. We perform a new installation of DB2 v9.7 64 bits for AIX.

Be sure to verify that your system meets all the minimum system requirements for installing the prerequisites, including adequate free disk space. See 4.2, "Preinstallation steps for AIX" on page 127

Before beginning the installation, be sure to log on to your system as a local system user with system rights authority.

Migration: In this section, we are dealing with a clean installation of Tivoli Storage Productivity Center V4.2. Therefore, be sure to understand that if you are required to migrate or upgrade your current Tivoli Storage Productivity Center environment to Tivoli Storage Productivity Center V4.2, you must follow the migration and upgrade sections found later in this chapter.

4.3.1 DB2 installation: Using the command line

This topic describes how to install DB2 v9.7 64-bit DB2 on AIX by using the command line. To install DB2, first log on as a user with root authority, and then use the following procedures.

Accessing the installation media by using the CD

Complete the following steps:

1. Create a mount point or choose an existing mount point. To create a mount point called /cdrom, enter the following command:

```
mkdir /cdrom
```

2. Insert the DB2 CD into the CD-ROM drive. Mount the CD-ROM file system at the desired mount point.

On AIX, you can use the **crfs** command to add an entry to /etc/file systems for the mount point. Run the following command:

```
/usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no' mount /cdrom
```

The **crfs** command has to be run only once for a given mount point, and after that you can use **mount** and **umount** for each CD or DVD you put in the drive.

3. Change to the directory where the CD-ROM is mounted:

```
cd /cdrom
```

Installing DB2

Complete the following steps:

1. At the command prompt on the host, execute the command line installer:

```
./db2_install
```

2. The installer is started, requesting you to either select the default installation directory, or optionally, select another directory, as shown in Figure 4-1. We choose **No**.

Figure 4-1 Select a directory

Select the product to install, ESE (DB2 Enterprise Server Edition) as shown in Figure 4-2.

```
**Specify one of the following keywords to install DB2 products.

ESE
CLIENT
RTCL

Enter "help" to redisplay product names.
Enter "quit" to exit.
```

Figure 4-2 Select product

Figure 4-3 shows the DB2 installation being initiated and informs you of the estimated time to perform all tasks.



Figure 4-3 DB2 ESE installation progress

The installation summary is displayed after a period of time and indicates a successful installation.

Configuring the DB2 environment

This section lists considerations for the users and groups that are required for the DB2 environment.

Complete the following steps:

1. Create users and groups for use with DB2 by typing the following commands from the host command line:

```
mkgroup id=999 db2iadm1
mkgroup id=998 db2fadm1
mkgroup id=997 dasadm1
mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1 home=/home/db2inst1 db2inst1
mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1 home=/home/db2fenc1 db2fenc1
mkuser id=1002 pgrp=dasadm1 groups=dasadm1 home=/home/dasusr1 dasusr1
```

2. Verify the owner of the directories by typing the 1s -1d for the directories (Figure 4-4); the directory owners are displayed as defined in step 1.

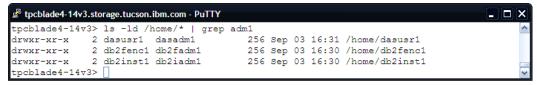


Figure 4-4 Verify directory owners

3. Set the DB2 user passwords, each of which you are required to enter twice:

```
passwd db2inst1
```

You are required to enter the password twice for verification, which presents the password that you want to use for the *DB2 instance*.

```
passwd db2fenc1
```

You are required to enter the password twice for verification, which presents the password that you want to use for the fe*nced user*.

```
passwd dasusr1
```

You are required to enter the password twice for verification, which presents the password that you want to use for the *DB2 administration server (DAS) user*.

4. Add authentication attributes to the users:

```
pwdadm -f NOCHECK db2inst1
pwdadm -f NOCHECK db2fenc1
pwdadm -f NOCHECK dasusr1
```

5. Change group db2iadm1 to include the root user:

chgroup users=db2inst1,root db2iadm1

Create a DB2 Administration Server (DAS):

/opt/IBM/db2/V9.7/instance/dascrt -u dasusr1

A message indicates that the program completed successfully. (Figure 4-5)



Figure 4-5 Create DAS server

7. Create a DB2 instance:

/opt/IBM/db2/V9.7/instance/db2icrt -a server -u db2fenc1 db2inst1

A message indicates that the program completed successfully (Figure 4-6).



Figure 4-6 Create DB2 instance

Source the instance profile:

. /home/db2inst1/sqllib/db2profile

Attention: There is a space between the period character (.) and /home.

8. Optional: Change the default location for database repositories. By default, this location is as follows:

/home/db2inst1

Space: The /home directory is usually not large enough for database repositories. Choose a file system with enough free space to contain the IBM Tivoli Storage Productivity repository.

In our case, we use the default repository location.

To change the default location, complete the following steps:

- a. Type the following command, where <new repository path> represents the new location for the repository:
 - db2 update dbm cfg using DFTDBPATH <new repository path> IMMEDIATE
- b. Type the following command, to assign ownership to db2inst1 and permission to anyone in db2iadm1:
 - chown -R db2inst1:db2iadm1 <new repository path>
- 9. Configure DB2 communication:
 - a. Edit the /etc/services file to verify or add the following line at the end of the file: db2c db2inst1 50000/tcp
 - b. Type the following commands:

```
db2 update dbm cfg using svcename db2c_db2inst1
db2set DB2COMM=tcpip
```

- 10.Add the DB2 license by typing the following two commands, where *<DB2 installer location>* represents the directory where the DB2 installer is located:
 - cd /opt/IBM/db2/V9.7/adm
 ./db2licm -a <DB2 installer location>/db2/ese/disk1/db2/license/db2ese o.lic
- 11. Restart DB2, by typing the following three commands, as shown in Figure 4-7:

```
db2stop force
db2 terminate
db2start
```

Figure 4-7 Restart DB2

4.3.2 Verifying that DB2 is installed correctly

The general steps to verify that DB2 has been installed properly are as follows:

- 1. Create the SAMPLE database.
- 2. Connect to the SAMPLE database.

- 3. Run a query against the SAMPLE database.
- 4. Drop the SAMPLE database.

To verify that DB2 has been installed successfully, complete the following steps:

 Change to the instance owner user ID by using the su command. For example, if your instance owner user ID is db2inst1, type the following command at the host command prompt:

```
su - db2inst1
```

- 2. Start the database manager by entering the **db2start** command.
- 3. Enter the db2samp1 command to create the SAMPLE database.
- 4. Enter the following DB2 commands from a DB2 command window to connect to the SAMPLE database, retrieve a list of all the employees that work in Department 20, and reset the database connection:

```
db2 connect to sample
db2 "select * from staff where dept=20"
db2 connect reset
```

The results of step 3 and step 4 are shown in Figure 4-8.

```
_ _ ×
🗗 tpcblade4-14v3.storage.tucson.ibm.com - PuTTY
  db2samp1
  Creating database "SAMPLE"...
  Connecting to database "SAMPLE"...
  Creating tables and data in schema "DB2INST1"...
  Creating tables with XML columns and XML data in schema "DB2INST1"...
  'db2sampl' processing complete.
$ db2 connect to sample
   Database Connection Information
Database server = DB2/AIX64 9.7.0
SQL authorization ID = DB2INST1
Local database alias = SAMPLE
$ db2 "select * from staff where dept=20"
       NAME
                 DEPT JOB YEARS SALARY COMM
 ---- ----- ----- -----
  10 Sanders 20 Mgr 7 98357.50 -
20 Pernal 20 Sales 8 78171.25 612.45
80 James 20 Clerk - 43504.60 128.20
190 Sneider 20 Clerk 8 34252.75 126.50
  4 record(s) selected.
$ db2 connect reset
DB20000I The SQL command completed successfully.
```

Figure 4-8 Verify DB2

5. If all steps completed successfully, you can remove the SAMPLE database by using the following command:

db2 drop database sample

4.4 Installing Tivoli Storage Productivity Center components

After the prerequisites are installed, we install the Tivoli Storage Productivity Center components.

4.4.1 Preparing for the installation

Before you begin the installation, consider the following requirements:

- ► Confirm that the correct version of DB2 is installed on your system.
- ▶ User IDs that will be required during the installation have been documented for reference.
- ▶ If you are planning to use LDAP, ensure that you have all the correct information.
- Make sure that DB2 is up and running.

We split the installation into two parts:

- 1. We install the Database Schema.
- 2. We install the remaining components, including Data server, Device server, Tivoli Integrated Portal, and Tivoli Storage Productivity Center for Replication.

Accessing the installation media using the CD

Complete the following steps:

1. Create a mount point or choose an existing mount point. To create a mount point called /cdrom, enter the following command:

```
mkdir /cdrom
```

2. Insert the CD into the CD-ROM drive. Mount the CD-ROM file system at the desired mount point.

On AIX, you can use the **crfs** command to add an entry to the /etc/filesystems location for the mount point. Run the following command:

```
/usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no' mount /cdrom
```

The **crfs** command has to be run only once for a given mount point, and after that you can use **mount** and **umount** for each CD or DVD that you put in the drive.

3. Change to the directory where the CD-ROM is mounted:

cd /cdrom

Accessing the installation media using a downloaded image

Complete the following steps:

- Create a temporary directory (for example, temp) to hold the Tivoli Storage Productivity Center installer tar files and untarred files. These files require 3 - 4 GB of hard drive space.
 mkdir /temp
- Copy or download the Tivoli Storage Productivity Center installer into temp directory.
- Change to the directory where you have stored the image, for example:cd /temp

4. Extract the Tivoli Storage Productivity Center installer file, by following the instructions supplied at the repository from which you download the image, which might involve running the tar or gunzip commands, or a combination of both, for example:

Be sure to extract both parts for disk1.

Change to the installation directory, which you extracted from the image, for example:cd /temp

Important: See 20.7, "Extracting the installation images on AIX" on page 796, for usage of the GNU tar program Version 1.14 or later, rather than the native AIX tar program, to extract files from Tivoli Storage Productivity Center installation images. The AIX tar program will truncate long file names, which can cause errors in Tivoli Storage Productivity Center and Tivoli Integrated Portal installation.

Preparing the display

If you are installing from a remote terminal session, you must set up an X Window System display or a Virtual Network Computing (VNC) Viewer connection prior to beginning the installation process.

If you decide to use X-Windows server, you first need to start your local X Window server application. Examples are Hummingbird Exceed, Cygwin or Xming. See Appendix C, "Configuring X11 forwarding" on page 815 for more information.

If you decide to use VNC Viewer, you must start the VNC server on the AIX server, set up a connection password, and then start the local VNC Viewer.

4.4.2 Creating the database schema

This topic provides information about how to create the database schema for use with Tivoli Storage Productivity Center.

Considerations:

- ▶ If you are using a remote database for Tivoli Storage Productivity Center, you must install the database schema on that computer after you have installed DB2.
- ► The DB2 database schema name for Tivoli Storage Productivity Center cannot be longer than eight characters.

Installing the database schema

Complete the following steps:

- Log on to the system with root authority.
- 2. Set up your shell environment to point to the instance where the database repository will be installed, to do this, source the **db2profile** script for the desired instance.

In our case, the DB2 instance is **db2inst1**, so we issue the following command:

. /home/db2inst1/sqllib/db2profile

Attention: There is a space between the period (.) and /home.

- 3. Change to the directory where you have extracted the Tivoli Storage Productivity Center software package, then launch the following command:
 - ./setup.sh
- 4. Tivoli Storage Productivity Center installer is launched, prompting you to select an installation language (Figure 4-9); click **OK** to continue.



Figure 4-9 Select language

5. Review the International License Agreement. If you agree with the terms, click I accept the terms of the license agreement, and then click Next (Figure 4-10).



Figure 4-10 License Agreement

Select **Custom installation** as the installation type (Figure 4-11). In addition, you can change the Tivoli Storage Productivity Center installation location to suit your requirements; we use the default location, which is /opt/IBM/TPC. Click **Next** to continue.



Figure 4-11 Custom Installation selection

6. In the next panel (Figure 4-12), remove all check marks except for **Create database schema** as specified during the DB2 installation. Click **Next**.

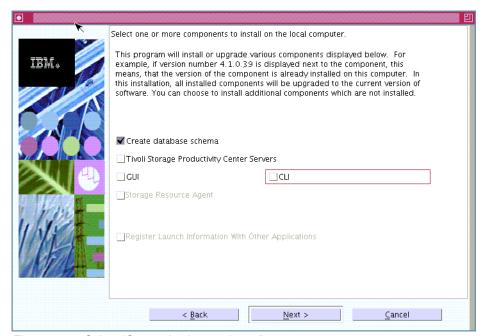


Figure 4-12 Select "Create database schema" component

7. The Database administrator information panel opens. Enter the user ID and password for the DB2 instance owner (Figure 4-13). Click **Next**.

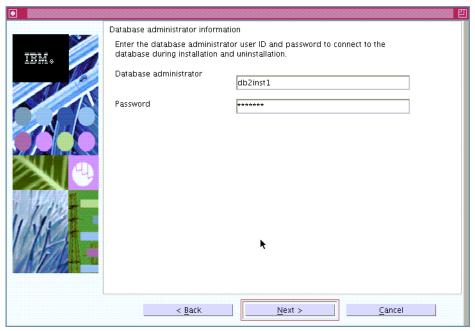


Figure 4-13 Database credentials

8. The new database schema information panel opens (Figure 4-14). Enter the DB user ID and password, and select **Create local database**.

If you click **Database creation details**, the Database schema creation information panel opens (Figure 4-15). Do not change the default values unless you are a knowledgeable DB2 administrator.

Click **Next** to continue.

See Appendix A, "DB2 table space considerations" on page 799 for the differences between SMS and DMS table spaces.



Figure 4-14 Database schema information

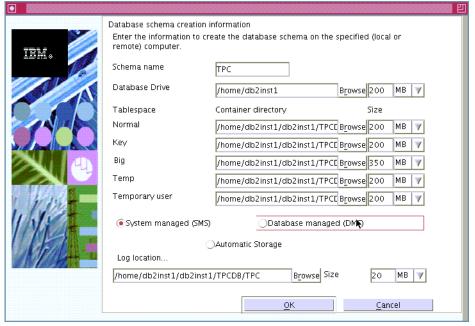


Figure 4-15 Database schema creation

9. The summary information panel opens (Figure 4-16). Click **Install** to begin the database schema installation.

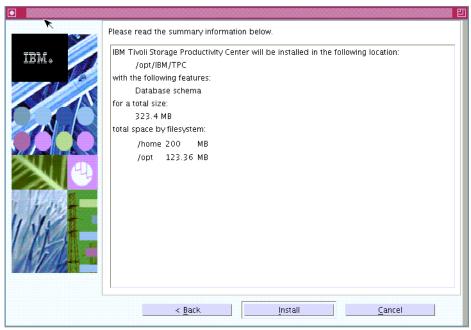


Figure 4-16 Summary information

The progress panel is displayed (Figure 4-17).

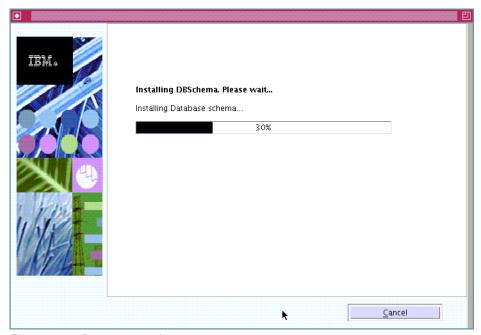


Figure 4-17 Progress panel

10. When the installation is complete, a summary of the installation is displayed (Figure 4-18). Click **Finish** to exit the installer.

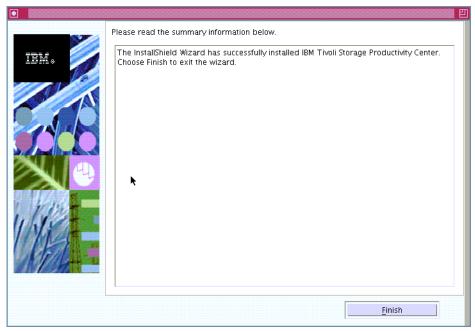


Figure 4-18 Installation results

Verifying database schema creation

To check the installation, verify that you have the database named TPCDB, as follows:

- 1. Source the db2 profile:
 - . /home/db2inst1/sqllib/db2profile
- 2. Verify the creation of the TPCDB database:

db2 list db directory

The command lists all existing databases, as shown in Figure 4-19.

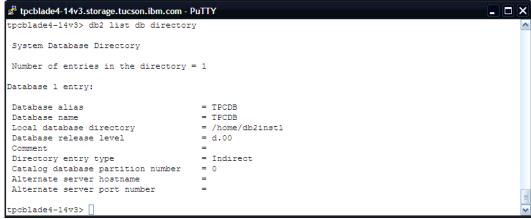


Figure 4-19 Verify database creation

4.4.3 Installing Tivoli Storage Productivity Center components

After completing the creation of the database schema, you are ready to install the following Tivoli Storage Productivity Center components:

- Data server
- Device server
- ► GUI
- ► CLI

In addition to these components, two additional components are automatically installed by default:

- Tivoli Integrated Portal
- Tivoli Storage Productivity Center for Replication

Use the following steps to complete the installation process:

- 1. Make sure that you are logged in with the root account.
- 2. Source the DB2 instance profile:
 - . /home/db2inst1/sqllib/db2profile
- 3. Change to the directory where you have extracted the Tivoli Storage Productivity Center software package, and then launch the following command:
 - ./setup.sh
- 4. Tivoli Storage Productivity Center installer is launched. Select an installation language (Figure 4-20) and click **OK** to continue.



Figure 4-20 Select language

5. The International License Agreement is displayed (Figure 4-21 on page 143). Review it and if you agree, click I accept the terms of the license agreement, and then click Next.



Figure 4-21 License Agreement

6. Select Custom installation type (Figure 4-22), and then click Next to continue.



Figure 4-22 Custom installation

- In the next panel (Select one or more components to install, Figure 4-23), select the following items, and then click Next:
 - Tivoli Storage Productivity Center Servers
 - GUI
 - CLI
 - Data agent (optional)
 - Fabric agent (optional)

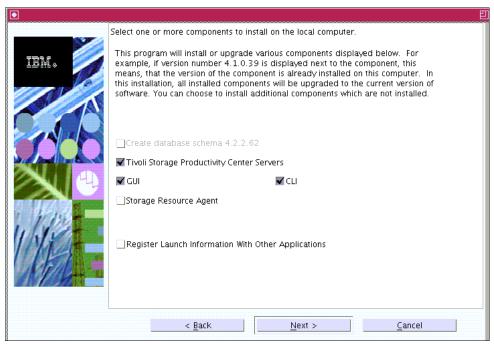


Figure 4-23 Select components

Attention: Preferably, *do not* install the Storage Resource agent at this time. Installing any Storage Resource agent by the installer requires you to also uninstall the Storage Resource agent by using the installer. Therefore, in most cases, doing this later using the Tivoli Storage Productivity Center GUI to deploy agents (instead of using the installer) is the more flexible approach.

8. If you are running the Tivoli Storage Productivity Center installation on a system with at least 4 GB but less than the preferable 8 GB of RAM, a warning message is issued (Figure 4-24). Ignore this message and continue with the installation by clicking **OK**.

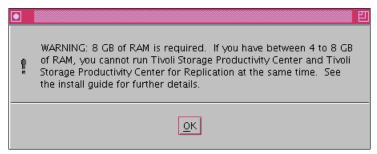


Figure 4-24 Memory size warning

Attention: When you attempt to install Tivoli Storage Productivity Center V4.2 on a system with less than 4 GB, an error message is issued and the installation fails.

The Database administrator information panel opens (Figure 4-25). The DB2 user ID and password are automatically filled in. The reason is because we used it to create the database schema. Click Next.

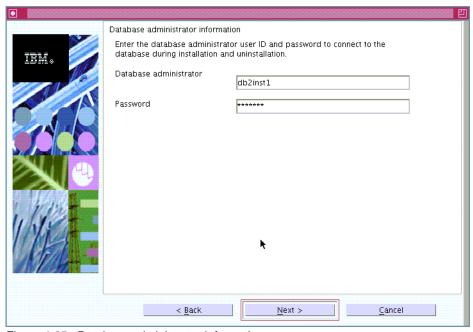


Figure 4-25 Database administrator information

10. The database schema panel opens (Figure 4-26). You have the option to select a local database or alternatively a remote database to be used by the Data server and Device server. We select the **Use local database**, because this is the database schema installed in the previous steps. Click **Next**.

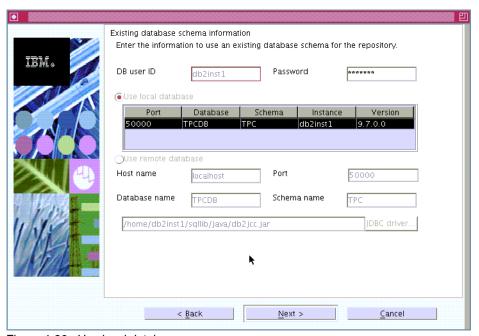


Figure 4-26 Use local database

11. The next panel, shown in Figure 4-27, requires the following input:

Important: Be sure to record all passwords that are used during the installation of Tivoli Storage Productivity Center.

Data Server Name:

Enter the fully-qualified host name of the Data server.

Data Server Port:

Enter the Data server port. The default is 9549.

- Device Server Name:

Enter the fully-qualified host name of the Device server.

- Device Server Port:

Enter the Device server port. The default is 9550.

- TPC superuser:

Enter an operating system group name to associate with the Tivoli Storage Productivity Center superuser role. This group must exist in your operating system before you install Tivoli Storage Productivity Center. Membership in this group provides full access to the Tivoli Storage Productivity Center product. You can assign a user ID to this group on your operating system and start the Tivoli Storage Productivity Center GUI using this user ID.

Override: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, the value you enter for the LDAP Tivoli Storage Productivity Center Administrator group overrides the value you entered here for the Tivoli Storage Productivity Center superuser.

Host authentication password:

This password is used by the Fabric agent to communicate with the Device server. This password must be specified when you install the Fabric agent.

- Data Server Account Password:

This password is not required for AIX installations; it is required only for Windows.

– WebSphere Application Server Admin ID and Password:

This WebSphere administrator user ID and password are required by the Device server to communicate with embedded WebSphere.

In our case, we use the db2inst1 user; you can use the Tivoli Storage Productivity Center superuser here. This user will be used for the local Tivoli Integrated Portal administrator ID.

Override: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the value you enter for the LDAP Tivoli Storage Productivity Center Administrator group overrides the value you entered here for the WebSphere Application Server admin ID and password. For LDAP considerations see Appendix D, "LDAP considerations" on page 833.

If you click Security roles, the Advanced security roles mapping panel opens. You can
assign a system group for each Tivoli Storage Productivity Center role that you want to
make an association with; this way allows you the flexibility to set up separate authority

IDs to perform various Tivoli Storage Productivity Center operations. The operating group must exist before you can associate a Tivoli Storage Productivity Center role with a group. You do not have to assign security roles at installation time; you can assign these roles after you have installed Tivoli Storage Productivity Center.

If you click **NAS discovery**, the NAS discovery information panel opens. You can enter the NAS filer login default user name and password and the SNMP communities to be used for NAS discovery. You do not have to assign the NAS discovery information at installation time, you can configure it after you have installed Tivoli Storage Productivity Center.Click Next to continue.

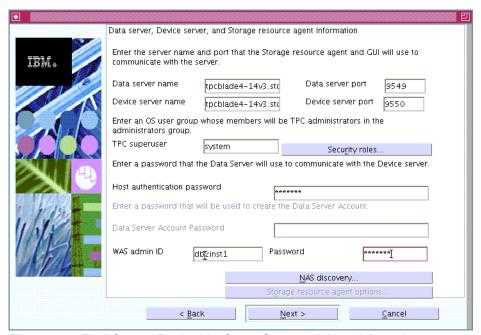


Figure 4-27 Tivoli Storage Productivity Center Server and Agent information

12. The Tivoli Integrated Portal (TIP) panel opens (Figure 4-28). You can select to install the Tivoli Integrated Portal program or use an existing Tivoli Integrated Portal install.

Important: Tivoli Integrated Portal must be installed on the same server as the Tivoli Storage Productivity Center server. You are limited to one Tivoli Storage Productivity Center instance for each Tivoli Integrated Portal.

Tivoli Integrated Portal will use 10 port numbers starting from the one specified in the Port field (referred to as the *Base Port*):

- base port+1
- base port+2
- base port+3
- base port+5
- base port+6
- base port+8
- base port+10 base port+12
- base port+13

The TIP Administrator ID and Password fields are prefilled with the WebSphere Application Server admin ID and password that are specified during step 11.

We choose to install the Tivoli Integrated Portal program and not use an existing Tivoli Integrated Portal. You have to specify the installation directory and the port to be used; we accept the defaults and click **Next** to continue.

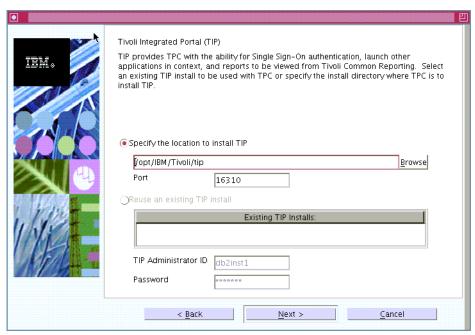


Figure 4-28 Tivoli Integrated Portal (TIP) panel

13. The authentication selection panel opens (Figure 4-29). This panel refers to the authentication method that will be used by Tivoli Storage Productivity Center to authenticate the users.

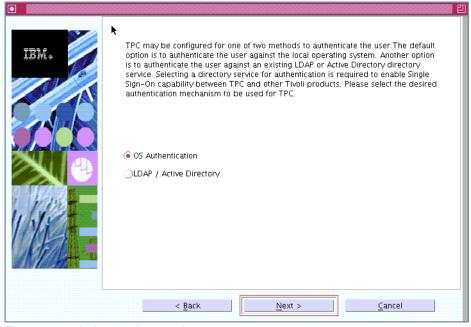


Figure 4-29 Authentication panel

If you have a valid Tivoli Integrated Portal instance on the system and it uses either OS-based or LDAP-based authentication, then Tivoli Storage Productivity Center will use that existing authentication method.

Otherwise, select the authentication method to use:

– OS Authentication:

This method uses the operating system for user authentication.

LDAP/Active Directory:

If you select LDAP or Microsoft Active Directory for authentication, you must have an LDAP or Active Directory already installed and set up.

Choose OS Authentication, and then click Next to continue.

14. The summary information panel is displayed (Figure 4-30). Review the information; at this stage, a good approach is to check that you have sufficient space in the required file systems. Click **Install** to continue.

Tip: Remember that the Replication Server is included in the installation of Tivoli Storage Productivity Center V4.2 by default, as mentioned before.

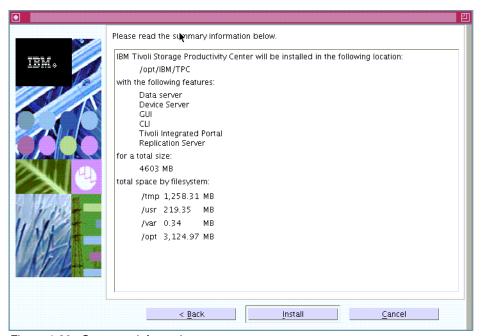


Figure 4-30 Summary information

15.A progress bar and progress message indicates the various stages within the installation process (Figure 4-31).

The installation starts with the Data server installation. The installer progresses through the separate components after a previous component is installed successfully.

Attention: If the installer fails to install a specific component, the process stops and the installer uninstalls all components.

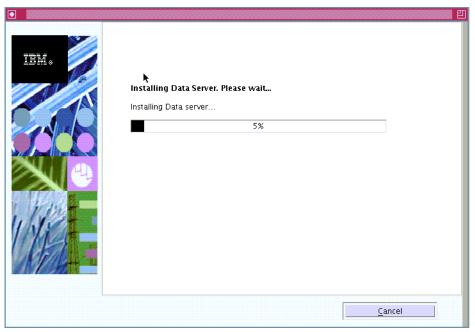


Figure 4-31 Data server install

The Device server is installed next (Figure 4-32). Various messages are displayed during the Device server installation process. When complete, the installer briefly displays the installation messages for the GUI, CLI, and Agents (if they were selected).

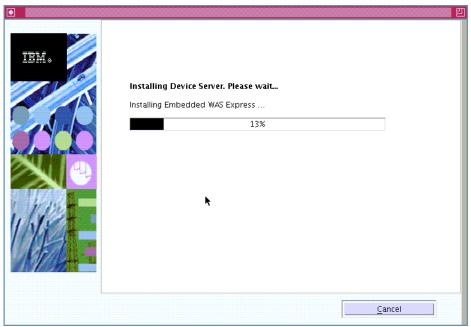


Figure 4-32 Device server installer

When done, the Installing TIP message is displayed (Figure 4-33).

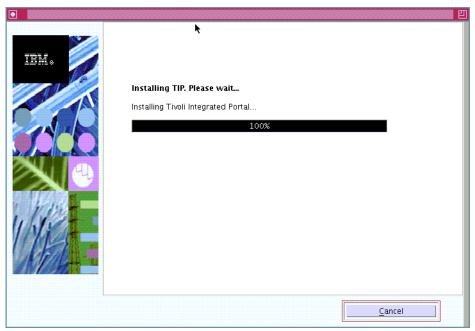


Figure 4-33 Tivoli Integrated Portal installing

Important: During the installation of Tivoli Integrated Portal on AIX systems, the progress bar incorrectly indicates that the Tivoli Integrated Portal installation is 100% complete even though it is not yet complete. *Continue to wait until the installation is complete.* The installation of Tivoli Integrated Portal can be a time consuming exercise, so be patient.

The Tivoli Storage Productivity Center installation is temporarily suspended. The *Tivoli Storage Productivity Center for Replication* is installed next.

Tivoli Storage Productivity Center for Replication installation

After the Tivoli Integrated Portal installation has completed, the Tivoli Storage Productivity Center for Replication installation is launched. The Tivoli Storage Productivity Center installation is temporarily suspended in the background, and the Tivoli Storage Productivity Center for Replication panel opens (Figure 4-34).



Figure 4-34 Tivoli Storage Productivity Center for Replication installation is launched

To install Tivoli Storage Productivity Center for Replication, complete the following steps:

1. On the Welcome panel (Figure 4-34), click **Next** to continue.

IMPORTANT: If you are not planning to use Tivoli Storage Productivity Center for Replication and you attempt to cancel or bypass the installation, an interruption in the installation process occurs, which invokes a complete Tivoli Storage Productivity Center installation rollback.

2. The System prerequisites check panel opens (Figure 4-35).

At this stage, the wizard checks that the operating system meets all prerequisite requirements and has the necessary fix packs installed.

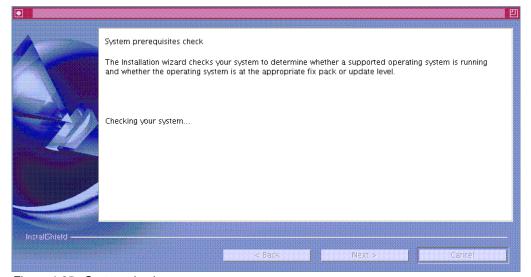


Figure 4-35 System check

3. If the system passes the check, as shown in Figure 4-36, click **Next** to continue.

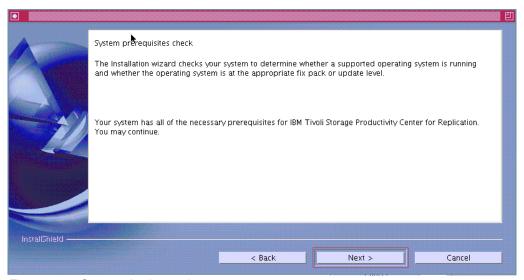


Figure 4-36 System check complete

4. Read the license agreement (Figure 4-37) and, if you agree, click Next to continue.

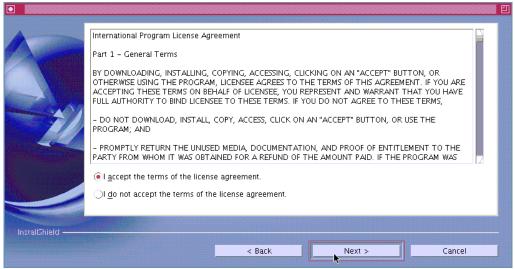


Figure 4-37 License agreement

 Select the Directory Name where you want to install Tivoli Storage Productivity Center for Replication. You can choose a directory either by changing the location or by accepting the default directory, as we have done in Figure 4-38. Click **Next** to continue.

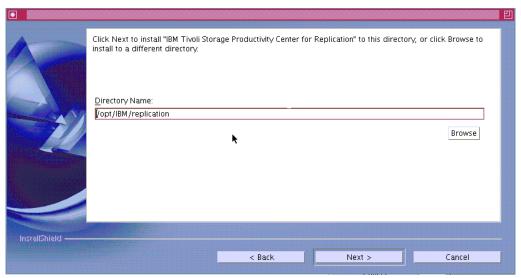


Figure 4-38 Directory Name

6. The Tivoli Storage Productivity Center Administrator user panel opens (Figure 4-39). Enter the user ID and password that will be used; this ID is usually the operating system administrator user ID. We choose the root user ID.

Tip: If you prefer to use another user ID, you must have already created it and ensure that it has administrator or system rights.

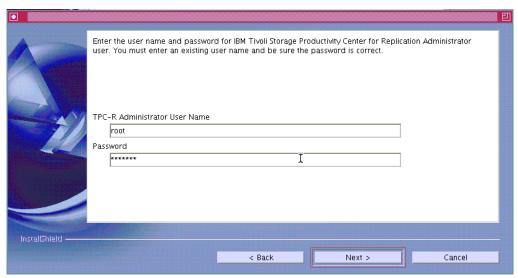


Figure 4-39 TPC-R User ID and Password

7. The Default WebSphere Application Server ports panel opens (Figure 4-40). Accept the defaults. Click **Next** to continue.

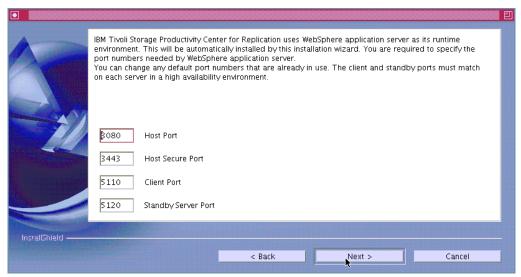


Figure 4-40 Default ports

8. The settings panel opens(Figure 4-41). Review the settings and make the necessary changes, if needed, by clicking **Back**. Otherwise, click **Install** to continue.

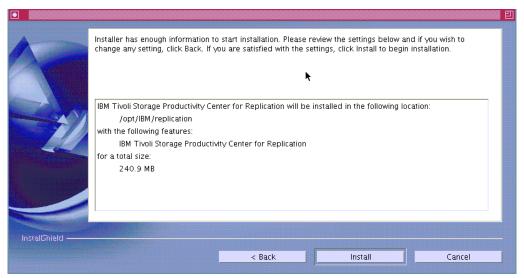


Figure 4-41 Summary display

The Tivoli Storage Productivity Center for Replication installation progress panel is displayed (Figure 4-42).

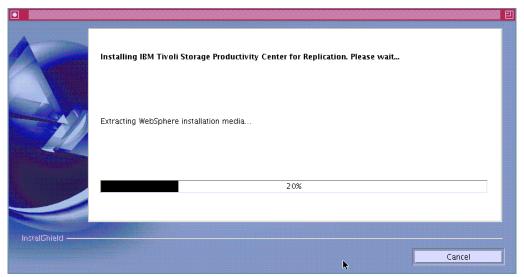


Figure 4-42 Tivoli Storage Productivity Center for Replication progress panel

9. A summary of the Tivoli Storage Productivity Center for Replication installation is displayed (Figure 4-43). Notice the URL that is listed for connecting to Tivoli Storage Productivity Center for Replication. Click **Finish**.

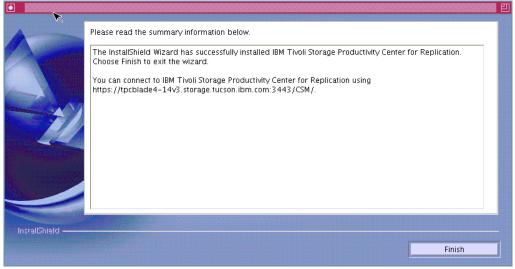


Figure 4-43 Installation results

Licenses: Tivoli Storage Productivity Center for Replication is installed with FlashCopy as the only licensed service. You must install the Two Site or Three Site Business Continuity (BC) license to use synchronous Metro Mirror and asynchronous Global Mirror capabilities.

You can now complete the installation of Tivoli Storage Productivity Center as described in the following topics.

Completing the installation of the Tivoli Storage Productivity Center

After the Tivoli Storage Productivity Center for Replication installation has completed, the Tivoli Storage Productivity Center Installer resumes, as shown in Figure 4-44.



Figure 4-44 Tivoli Storage Productivity Center installation continues

The Tivoli Storage Productivity Center installation results panel is displayed (see Figure 4-45). Click **Finish** to continue.

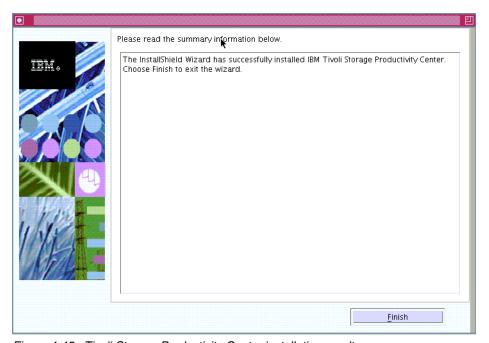


Figure 4-45 Tivoli Storage Productivity Center installation results

Verifying the installation

At the end of the installation, a good approach is to make sure that all the components have been installed successfully and that Tivoli Storage Productivity Center is in good working order.

To test this on AIX, we launch **Tivoli Integrated Portal** → **Tivoli Storage Productivity Center** user interface. On Tivoli Storage Productivity Center, we confirm that all servers are started and running:

- 1. Launch the Tivoli Integrated Portal using the URL that is specific to your environment: https://tpc_server_name:16316/ibm/console/logon.jsp
- 2. Log in using the root account, as shown in Figure 4-46.

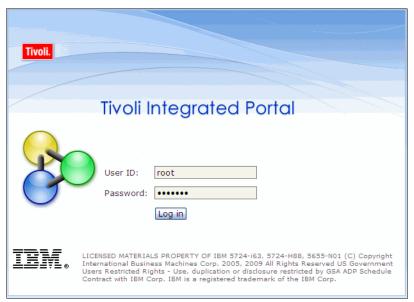


Figure 4-46 Tivoli Integrated Portal Login

3. Start the Tivoli Storage Productivity Center user interface (Figure 4-47).



Figure 4-47 Tivoli Storage Productivity Center user interface

4. Verify that all services are started (Figure 4-48), the nodes are reflected as green.

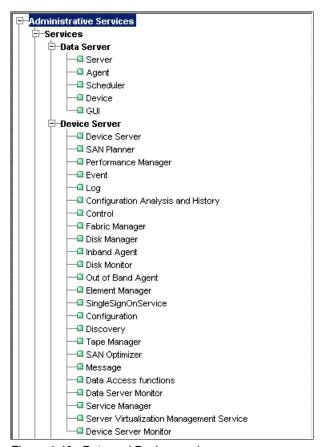


Figure 4-48 Data and Device services



Migrating Tivoli Storage Productivity Center base code to current level

In this chapter, we describe migration of the Tivoli Storage Productivity Center base code to Tivoli Storage Productivity Center V4.2 level. We explain special considerations that you have to be aware of during the upgrade.

This chapter contains the following topics:

- Migration considerations and scenarios
- Prerequisites for upgrade to V4.2
- Database considerations
- ► Tivoli Storage Productivity Center for Replication considerations
- Agent Manager, Data agent, and Fabric agent considerations
- Storage Resource agent migration
- ▶ Upgrading Tivoli Storage Productivity Center for Replication in high availability relationship
- Upgrading System Storage Productivity Center (SSPC)

5.1 Migration considerations

If you are planning to upgrade your existing Tivoli Storage Productivity Center to Tivoli Storage Productivity Center V4.2, you have to plan migration steps and consider prerequisites before you do the installation. To upgrade Tivoli Storage Productivity Center, use the same installation program as you used for installing the product. When you upgrade Tivoli Storage Productivity Center, you are upgrading all installed components including the database schema. If a component fails to upgrade, then only that component is not upgraded.

You can migrate previous Tivoli Storage Productivity Center V3.3.2 or later releases and Tivoli Storage Productivity Center for Replication V3.3.2 to Tivoli Storage Productivity Center V4.2.

Migration:

- ► TotalStorage Productivity Center (TPC) was renamed Tivoli Storage Productivity Center in Version 4.1.
- ► A best practice is to migrate from Tivoli Storage Productivity Center V3.3.2 to V4.1, and then upgrade from V4.1 to V4.2.

Tivoli Storage Productivity Center for Replication is no longer a stand-alone application (since Tivoli Storage Productivity Center V4.1). Therefore, when you upgrade from Tivoli Storage Productivity Center for Replication V3.3.2, Tivoli Storage Productivity Center Version V4.2 will install Tivoli Integrated Portal and Tivoli Storage Productivity Center for Replication V4.2.

5.1.1 Prerequisites

Before starting the upgrade, be sure that your system meets the hardware and software requirements of Tivoli Storage Productivity Center V4.2. Verify the requirements at the following location:

http://www.ibm.com/support/entry/portal/Planning/Software/Tivoli/Tivoli_Storage_Pr
oductivity_Center_Standard_Edition

5.1.2 Database considerations

If you are planing to upgrade Tivoli Storage Productivity Center, you must consider the database repository because Tivoli Storage Productivity Center V4.2 supports only DB2 as the database repository.

The following DB2 versions are supported with Tivoli Storage Productivity Center V4.2:

- V9.1 (Fix Pack 2 or later, not supported on Windows 2008 R2. Use DB2 9.5 or 9.7)
- ► V9.5 (Fix Pack 6 not supported)
- ► V9.7 (Fix Pack 1 to 3a not supported)

For the most current Tivoli Storage Productivity Center 4.2 platform support and database repository support, see the following website:

https://www-304.ibm.com/support/docview.wss?uid=swg27019380

If you have DB2 Version 8.2, DB2 Version 9.1, or DB2 Version 9.5 installed, migrate and upgrade to DB2 Version 9.7.

The general steps are as follows:

- 1. Stop the IBM Tivoli Storage Productivity Center services and Agent Manager (if you have Agent Manager installed).
- 2. Verify the database migration path.
- 3. Back up the database.
- Install DB2 9.7
- 5. Migrate the DB2 instance.
- 6. Migrate the database.
- 7. Verify the migration.
- 8. Start the IBM Tivoli Storage Productivity Center services and Agent Manager (if you have Agent Manager installed).

For more information about the upgrade to DB2 Version 9.7, see the following resources:

- ▶ IBM Tivoli Storage Productivity Center Installation and Configuration Guide, SC27-2337
- ► For more information about the upgrade to DB2 Version 9.7, see "Upgrade to DB2 Version 9.7" in the IBM DB2 Information Center at the following location:

http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.qb.upgrade.doc/doc/c0023662.html

Tip: If the Tivoli Storage Productivity Center database is on a remote system from the server, you must also upgrade the remote database.

5.1.3 Tivoli Storage Productivity Center for Replication considerations

To upgrade to V4.2, uninstalling the previous version of Tivoli Storage Productivity Center for Replication is unnecessary. Tivoli Storage Productivity Center V4.2 can be installed on an existing V3.x or 4.x installation if the hardware and software requirements needed to support Tivoli Storage Productivity Center are met.

With V4.2, Tivoli Storage Productivity Center for Replication no longer supports DB2 as the data store for its operational data. Tivoli Storage Productivity Center for Replication uses an embedded repository (Derby database) for its operational data.

The Tivoli Storage Productivity Center for Replication V4.2 installation program automatically migrates any Tivoli Storage Productivity Center for Replication operational data in an existing Tivoli Storage Productivity Center for Replication DB2 database to the Tivoli Storage Productivity Center for Replication embedded repository as part of upgrading to Tivoli Storage Productivity Center for Replication V4.2 from an earlier version.

If you do not use or plan to use Tivoli Storage Productivity Center for Replication, do not interrupt the upgrade installation by clicking the **Cancel** button on InstallShield Wizard for Tivoli Storage Productivity Center for Replication, this will result in a failed upgrade. Click **Next** and finish the installation of Tivoli Storage Productivity Center for Replication.

Tip: If you already have Tivoli Storage Productivity Center for Replication disabled, you do not have to start it before Tivoli Storage Productivity Center upgrade. InstallShield Wizard will start the Tivoli Storage Productivity Center for Replication service and continue with the installation.

If you do not plan to use Tivoli Storage Productivity Center for Replication, you can disable it after the upgrade.

Disabling Tivoli Storage Productivity Center for Replication

This section describes the steps for Windows, Linux, and AIX.

On Windows

Complete the following steps;

- To disable the Tivoli Storage Productivity Center for Replication server, click Start → Settings → Control Panel → Administrative Tools → Services.
- 2. Right-click the following service and then click **Properties**, as shown in Figure 5-1: IBM WebSphere Application Server V6.1 CSM

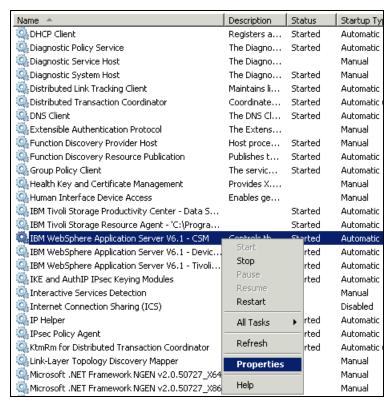


Figure 5-1 Tivoli Storage Productivity Center for Replication server service properties

3. On the next panel (Figure 5-2), select **Disabled** under the Startup type menu and then click **Stop**. After the service is stopped, click **OK** to close this panel.

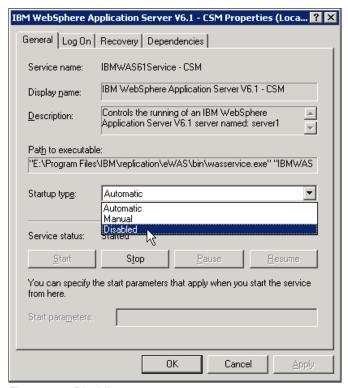


Figure 5-2 Disabling server

On Linux and AIX

Complete the following steps:

1. To stop the Tivoli Storage Productivity Center for Replication Server on Linux and AIX issue the following command from the command prompt. See Figure 5-3:

/opt/IBM/replication/eWAS/profiles/CSM/bin/stopServer.sh server1 -username
<username> -password <password>

In the command, <username> is the user ID, and <password> is the password created during installation.

```
| root@baltic:/opt/ibm | ropt/IBM/replication/eWAS/profiles/CSM/bin/stopServer.sh server1 -username dasusr1 -pa sword dasusr1 | ADMU01161: Tool information is being logged in file /opt/IBM/replication/eWAS/profiles/CSM/logs/server1/stopServer.log | ADMU01281: Starting tool with the CSM profile | ADMU31001: Reading configuration for server: server1 | ADMU31001: Rever stop request issued. Waiting for stop status. | ADMU40001: Server server1 stop completed. | Iroot@baltic ibm]# |
```

Figure 5-3 Stop Tivoli Storage Productivity Center for Replication server

2. To *disable* the Tivoli Storage Productivity Center for Replication Server from starting on system reboot, edit the /etc/inittab file by adding the number sign character (#) to the beginning of the line that starts the *Tivoli Storage Productivity Center for Replication*, as shown in Figure 5-4.

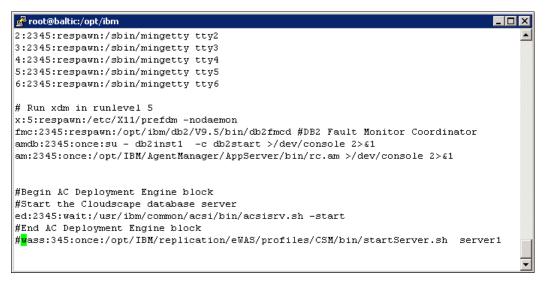


Figure 5-4 Edit /etc/inittab

If you plan to use only Tivoli Storage Productivity Center for Replication, you can disable Tivoli Storage Productivity Center after the upgrade.

Disabling Tivoli Storage Productivity Center

This section describes the steps for Windows, Linux, and AIX.

On Windows

Complete the following steps:

To disable Tivoli Storage Productivity Center, click Start → Settings → Control Panel → Administrative Tools → Services.

2. Right-click the following service and then select **Properties** (Figure 5-5): IBM WebSphere Application Server V6.1 - DeviceServer

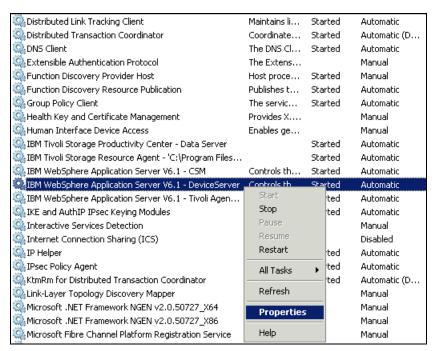


Figure 5-5 Service properties

3. On the next panel (Figure 5-6), select **Disabled** under the Startup type menu, and then click **Stop**. After the service is stopped, click **OK** to close this panel.

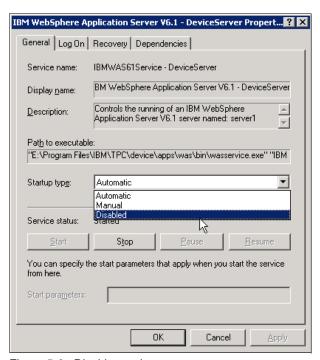


Figure 5-6 Disable service

- 4. Repeat the same procedure for the following services:
 - IBM Tivoli Storage Productivity Center Data Server
 - IBM Tivoli Common Agent <directory>

(<directory> is where the Common Agent is installed. The default is
<TPC_install_directory>\ca)

IBM Tivoli Storage Resource agent - <directory>

(<directory> is where the Storage Resource agent is installed. The default is
<TPC_install_directory>\agent)

- Tivoli Integrated Portal TIPProfile_Port_<xxxxx>
 (<xxxxx> indicates the port specified during installation. The default port is 16310.)
- IBM ADE Service (Tivoli Integrated Portal registry)

Attention: Stop Tivoli Integrated Portal and IBM ADE Service only if no other applications are using these services.

On Linux

Complete the following steps:

- 1. To stop the Tivoli Storage Productivity Center services, as shown in Figure 5-7, run the following commands in the command prompt window:
 - Data server:

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 stop

Device server:

/<usr or opt>/IBM/TPC/device/bin/linux/stopTPCF.sh

- 2. Depending on whether you have a Data agent or Storage Resource agent installed, issue the following commands accordingly:
 - Common Agent:

/<usr or opt>/IBM/TPC/ca/endpoint.sh stop

Storage Resource agent:

/<usr or opt>/IBM/TPC/agent/bin/agent.sh stop

Figure 5-7 Stop Tivoli Storage Productivity Center services on Linux

On AIX

Complete the following steps:

- To stop the Tivoli Storage Productivity Center services as seen in Figure 5-7 on page 168, run the following commands in the command prompt window:
 - Data server:

```
stopsrc -s TSRMsrv1
```

Device server:

```
/<usr or opt>/IBM/TPC/device/bin/aix/stopTPCF.sh
```

- Depending on whether or not you have a Data agent or Storage Resource agent installed, issue the following commands:
 - Common Agent:

```
/<usr or opt>/IBM/TPC/ca/endpoint.sh stop
```

Storage Resource agent:

```
/<usr or opt>/IBM/TPC/agent/bin/agent.sh stop
```

3. To *disable* the Tivoli Storage Productivity Center Server from starting when the system reboots, you must edit the /etc/inittab by adding a number sign character (#) at the beginning of the line that starts the *Tivoli Storage Productivity Center*, as shown in Figure 5-8.

```
💤 9.12.5.35 - PuTTY
                                                                            logsymp:2:once:/usr/lib/ras/logsymptom # for system dumps
perfstat:2:once:/usr/lib/perf/libperfstat updt dictionary >/dev/console 2>&1
diagd:2:once:/usr/lpp/diagnostics/bin/diagd >/dev/console 2>&1
rcwpars:2:once:/etc/rc.wpars > /dev/console 2>&1 # Corrals autostart
xmdaily:2:once:/usr/bin/xmwlm -L 2>&1 >/dev/null # Start xmwlm daily recording
ctrmc:2:once:/usr/bin/startsrc -s ctrmc > /dev/console 2>&1
dt:2:wait:/etc/rc.dt
cons:0123456789:respawn:/usr/sbin/getty/dev/console
pconsole:2:once:/usr/bin/startsrc -s pconsole > /dev/null 2>&1
ha star:h2:once:/etc/rc.ha star >/dev/console 2>&1
tty0:2:off:/usr/sbin/getty/dev/tty0
fmc:2:respawn:/opt/IBM/db2/V9.5/bin/db2fmcd #DB2 Fault Monitor Coordinator
amdb:2345:once:su - db2inst1 -c db2start >/dev/console 2>&1
am:2345:once:/opt/IBM/AgentManager/AppServer/bin/rc.am >/dev/console 2>&1
#TSRMsrv1:2:once:/usr/bin/startsrc -s TSRMsrv1 >/dev/console 2>&1
:Begin AC Deployment Engine block
:Start the Cloudscape database server
si:23456789:wait:/usr/ibm/common/acsi/bin/acsisrv.sh -start
:End AC Deployment Engine block
wass:2:once:/opt/IBM/replication/eWAS/profiles/CSM/bin/startServer.sh server1
```

Figure 5-8 Disable Tivoli Storage Productivity Center

Stopping Tivoli Integrated Portal on AIX and Linux

Complete the following steps:

 To stop Tivoli Integrated Portal, run the following command in a command prompt window, as shown in Figure 5-9:

```
<install_directory>/tip/profiles/TIPProfile/bin/stopServer server1
-username <tipadmin>
-password <password>
```

In the command, <tipadmin> is the administrator user ID, and <password> is the administrator password. Wait for the server to complete the operation.

- 2. To stop the IBM ADE Service, run the following commands in a command prompt window:
 - Source the environment: . /var/ibm/common/acsi/setenv.sh
 Run the command: /usr/ibm/common/acsi/bin/acsisrv.sh stop

- Hun the command: /usr/lbm/common/acs1/bln/acs1srv.sn stop

Attention: Stop Tivoli Integrated Portal and IBM ADE Service only if no other applications are using these services.

```
root@baltic:/opt/IBM/Tivoli/tip/profiles/TIPProfile/bin

[root@baltic bin] # /opt/IBM/Tivoli/tip/profiles/TIPProfile/bin/stopServer.sh server1 -
-username root -password itso13sj

ADMU0116I: Tool information is being logged in file
/opt/IBM/Tivoli/tip/profiles/TIPProfile/logs/server1/stopServer.log

ADMU0128I: Starting tool with the TIPProfile profile
ADMU3100I: Reading configuration for server: server1

ADMU3201I: Server stop request issued. Waiting for stop status.

ADMU4000I: Server server1 stop completed.

[root@baltic bin] #
[root@baltic bin] #
[root@baltic bin] #
```

Figure 5-9 Stopping Tivoli Integrated Portal

5.2 Credentials migration tool

With Tivoli Storage Productivity Center V4.2, new native storage system interfaces are provided for DS8000, SAN Volume Controller, and XIV storage systems. The native interfaces replace the CIM agent (SMI-S agent) implementation for these storage systems. By upgrading Tivoli Storage Productivity Center to V4.2, you must migrate the existing storage system credentials for the native interfaces.

If you are upgrading Tivoli Storage Productivity Center, a storage subsystem credentials migration tool is provided to help you migrate the existing storage system credentials for the native interfaces. The migration tool is able to migrate the existing storage system credentials automatically after you confirm the migration of the credentials of devices by using the update button in the migration tool.

The native interfaces are supported for the following release levels:

- ► DS8000: release 2.4.2 or later
- ► SAN Volume Controller (SVC): Version 4.2 or later
- ► XIV: Version 10.1 or later

If Tivoli Storage Productivity Center can access the DS8000 by using the existing credentials, you can continue to use the existing credentials. For XIV and SAN Volume Controller storage systems, you must manually update the credentials.

5.2.1 Migration methods

You can migrate your storage system credentials in one of the following ways:

Before:

Use the migration tool before the upgrade from the Tivoli Storage Productivity Center installation packages.

► During:

Run the migration tool during the Tivoli Storage Productivity Center upgrade procedure. Run the Tivoli Storage Productivity Center upgrade installation program and specify that you want to use the migration tool.

After:

Use the Tivoli Storage Productivity Center GUI after the upgrade to migrate the credentials.

Considerations:

- Use the migration tool before you start upgrade to Tivoli Storage Productivity Center V4.2 to check and prepare all your monitored and managed devices to be ready immediately after the upgrade. The migration tool will store the credentials into the database repository.
- If a device is in the probe definition before the upgrade to Tivoli Storage Productivity Center V4.2, it will not show up in the Configuration Device wizard because the Configuration Device wizard is only for configuring those devices that are unconfigured for monitoring. In this case, run the migration tool before the upgrade.

5.2.2 Running the credentials migration tool before the upgrade

You can migrate your storage system credentials before you start the upgrade of Tivoli Storage Productivity Center. Run the migration tool from the UserMigrationTool directory by starting MigrateUserInfo.bat (for Windows) or MigrateUserInfo.sh (for UNIX or Linux). It opens the User Credentials Migration Tool window, which lists the subsystems that can be updated (Figure 5-10).

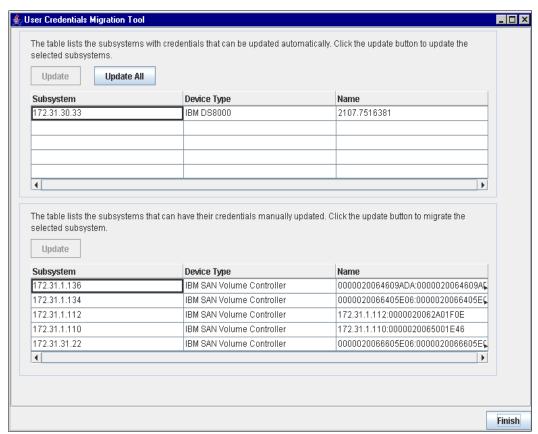


Figure 5-10 User Credentials Migration Tool window

If you run the User Credentials Migration Tool after the upgrade to Tivoli Storage Productivity Center V4.2, an error occurs (Figure 5-11). In this case, you have to run the User Credentials Migration Tool from the Tivoli Storage Productivity Center GUI.



Figure 5-11 User Credentials Migration Tool error

Tip: To run the tool on Windows, a new DLL, msvcr90.d11, is required. If it is not installed, the migration tool cannot start. If that happens, start Tivoli Storage Productivity Center 4.2 installer, choose the language, and accept the license terms. At that point, the required DLL will be installed. You can go back and launch the stand-alone migration tool.

5.2.3 Running the credentials migration tool during the upgrade

You can migrate your storage system credentials during the upgrade of Tivoli Storage Productivity Center. After you start the Tivoli Storage Productivity Center installer program and after you specify database repository information in the installer window, a window opens (Figure 5-12) where you specify that you want to run Storage Subsystem Credential Migration Tool.

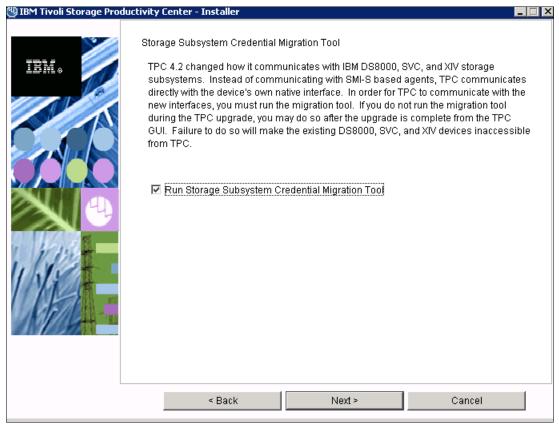


Figure 5-12 Credential Migration Tool selection within Tivoli Storage Productivity Center GUI installer

If you select that you want to run the tool, the User Credentials Migration Tool window opens after you click **Install** button on the summary window. The window shows the table list of the subsystems that can be updated (Figure 5-10 on page 172).

5.2.4 Running the credentials migration tool after the upgrade

You can migrate your storage system credentials after you successfully upgrade Tivoli Storage Productivity Center to V4.2. When you start and log into Tivoli Storage Productivity Center GUI, the welcome window shown in Figure 5-13 opens.

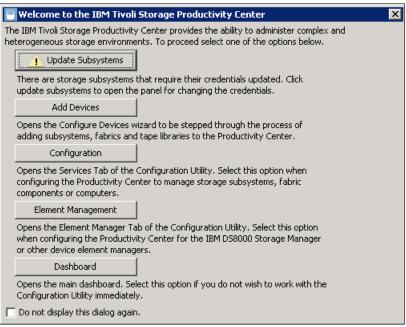


Figure 5-13 Tivoli Storage Productivity Center GUI Welcome window

Click **Update Subsystems**. The panel for changing credentials opens (Figure 5-14).

If you close the welcome window, you can also open the panel from navigation tree by clicking $Configuration \rightarrow Update Storage Subsystem Credentials.$

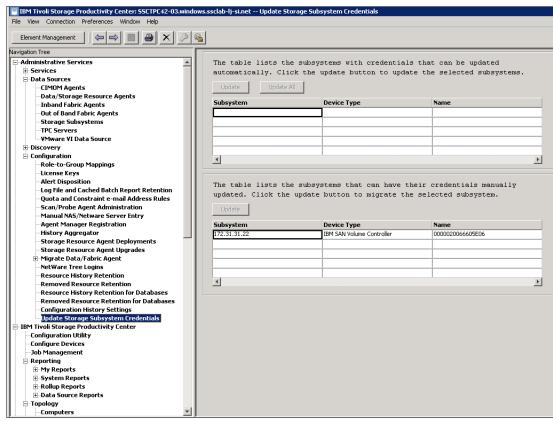


Figure 5-14 Update Storage Subsystem Credentials panel

Tip: The storage system credential migration applies to all DS8000 systems, XIV systems, and SAN Volume Controller systems. If you have run a CIMOM discovery job for a storage system but have not run a probe job for that system, then upgrade Tivoli Storage Productivity Center, the IP address does not display in the GUI. You must manually enter the IP address for that storage system.

5.3 Agent Manager, Data agent, Fabric agent considerations

Because the Storage Resource agents now perform the functions of the Data agents and Fabric agents, you no longer need to install and maintain the Agent Manager, Data agents, and Fabric agents. The Data and Fabric agents are no longer part of Tivoli Storage Productivity Center and they are not included with Tivoli Storage Productivity Center V4.2. The Agent Manager is included with Tivoli Storage Productivity Center V4.2 in case clients need to upgrade it to work with DB2 9.7 in order to continue using legacy agents. The Tivoli Storage Productivity Center V4.2 supports existing Data and Fabric agents and Agent Manager but without any new functions and with some limitations.

By migrating Tivoli Storage Productivity Center to V4.2, Data and Fabric agents can be migrated using the migration function that has been developed to assist you with migrating to the new Storage Resource agents. The migration function was designed so that if the migration fails, the Data and Fabric agents are restored and restarted. You can either continue to use the Data and Fabric agents or retry the migration. If you choose not to migrate the legacy agents as part of the server upgrade, the graphical installer can be launched at a later time to migrate the legacy agents.

Data and Fabric agent migration and upgrade is described in detail in 5.3, "Agent Manager, Data agent, Fabric agent considerations" on page 175.

Considerations:

- ► Any version of Agent Manager 1.3.2 supports DB 9.1. For DB2 9.5 support, you need to use Agent Manager Version 1.3.2.30, which is included with Tivoli Storage Productivity Center 4.1.1.
- ► If you are planning to use DB2 9.7, you must install Agent Manager 1.4.x or later. Agent Manager 1.3.x does not support DB2 v9.7.

The Tivoli Storage Productivity Center V4.2 installation program does not support installation of the legacy Data agent or Fabric agent. If you want to install the legacy Data agent or Fabric agent, you must have a previous Tivoli Storage Productivity Center installation program that supports installing the Data agent or Fabric agent.

If you plan to use existing legacy Data and Fabric agents, Tivoli Storage Productivity Center must be registered to Agent Manager. By migrating Tivoli Storage Productivity Center base code to V4.2, an existing Agent Manager will remain registered. If plan to migrate Data and Fabric agents to the new Storage Resource agents, Agent Manager is no longer required and it can be uninstalled.

Important: When running with a Tivoli Storage Productivity Center V4.2 server and a Data agent Version 3.3.x or 4.1.x, you see the following limitations:

- ▶ When you are using a Tivoli Storage Productivity Center V4.2 server and a Data agent lower than Version 4.1.0, you get error messages in the logs for the storage subsystem performance and switch performance reports (GEN0324E and GEN0008E) if there is data. These error messages do not affect the reports. The report job ends with a warning message. The job status is correct, and the job log reflects the results of the report.
- ► The performance constraint violation reports will not be able to run with a Tivoli Storage Productivity Center V4.2 server and a Data agent Version 4.1.0 or lower. The Data agents have been removed from the agent list. You can migrate the Data agent to a Storage Resource agent to get a performance constraint violation report.
- You cannot create a batch report for Rollup Reports by clicking IBM Tivoli Storage Productivity Center → Reporting → Rollup Reports → Asset → Computers → By Computer. The Data agents have been removed from the agent list. You can migrate the Data agent to a Storage Resource agent to get a batch report for Rollup Reports.

If you have Tivoli Storage Productivity Center V4.1.1 (or earlier) agents installed, and if you want to continue to use them, Table 5-1 shows the valid upgrade scenarios.

Table 5-1 Agent upgrade scenarios

Tivoli Storage Productivity Center agent V4.1.1 (or earlier) installed	Use Tivoli Storage Productivity Center V4.1.1 installation program on non-Tivoli Storage Productivity Center server for local installation
Data agent or Fabric agent or both installed (Version 4.1.1 or earlier) on local machine	If the Data agent or Fabric agent is down level, the agent will be upgraded to the latest V4.1.1 level. If the Data agent or Fabric agent is at the latest V4.1.1 level, you see a message that the agent is already installed.
Storage Resource agent is installed on local machine	If the Storage Resource agent is at the latest V4.1.1 level, the Storage Resource agent is left as is. If the Storage Resource agent is not at the latest V4.1.1 level, the agent is migrated to a Data agent or Fabric agent.
No agent installed	The Data agent or Fabric agent is installed.

5.4 Migration scenarios

Depending on your existing installation, the upgrade migration scenarios, listed in Table 5-2, are possible.

Table 5-2 Agent migration scenarios

Tivoli Storage Productivity Center 4.1.1 or earlier versions	Migration scenarios to Tivoli Storage Productivity Center 4.2
Data agent or Fabric agent or both are installed	You have a choice: ► Leave the Data agent or Fabric agent at the down level version ► Migrate the Data agent or Fabric agent to Storage Resource agent
Storage Resource agent is installed	You have a choice to upgrade or not upgrade the Storage Resource agent to 4.2.
No Agent installed	The default Storage Resource agent is installed

Considerations:

- ▶ You cannot use the Tivoli Storage Productivity Center V4.1.1 (or earlier) installation program on a Tivoli Storage Productivity Center V4.2 system.
- You can use the Tivoli Storage Productivity Center V4.2 installation program to install a local Storage Resource agent on a system that does not have the Tivoli Storage Productivity Center server installed. You can also use the Tivoli Storage Productivity Center GUI to deploy the Storage Resource agents (from the server system).

5.4.1 Migration from Tivoli Storage Productivity Center V3.x

You can upgrade previous Tivoli Storage Productivity Center V3.3.2, or later releases, and Tivoli Storage Productivity Center for Replication V3.3.2 to Tivoli Storage Productivity Center V4.2.

The best practice is to upgrade in two steps:

- ► Step 1:
 - Upgrade DB2 to Version 9.1.
 - Upgrade from Tivoli Storage Productivity Center V3.3.2 to Tivoli Storage Productivity Center V4.1.
- ► Step 2:
 - Upgrade DB2 to Version 9.7.
 - Upgrade from Tivoli Storage Productivity Center V4.1 to Tivoli Storage Productivity Center V4.2.

Considerations:

- ► TotalStorage Productivity Center (TPC) was renamed Tivoli Storage Productivity Center in Version 4.1.
- ► Upgrading from Tivoli Storage Productivity Center Version 3.3.2 to Tivoli Storage Productivity Center Version 4.1 is described in *IBM Tivoli Storage Productivity Center V4.1 Release Guide*, SG24-7725, Chapter 3.

5.4.2 Migration from Tivoli Storage Productivity Center V4.1

You can directly upgrade existing Tivoli Storage Productivity Center V4.1 and Tivoli Storage Productivity Center for Replication V4.1 instances to Tivoli Storage Productivity Center V4.2. Verify that your existing database repository is supported with Tivoli Storage Productivity Center V4.2.

In this section, we show you how to upgrade the Tivoli Storage Productivity Center components when you have installed Data and Fabric agents. Chapter 6, "Agent migration and upgrade" on page 209 shows you how to upgrade the agents after the Tivoli Storage Productivity Center is successfully upgraded to V4.2

Before proceeding with the upgrade, several steps must be completed.

Preparing for Tivoli Storage Productivity Center component migration Complete the following steps:

- 1. Exit all instances of the Tivoli Storage Productivity Center GUI.
- 2. If you are upgrading Tivoli Storage Productivity Center on a Windows server, make sure that you have exclusive access to the server you are installing Tivoli Storage Productivity Center V4.2 on. If you are accessing the server remotely, make sure that there are no other remote connections to the server. Multiple remote connections, such as Windows Remote Desktop Connections, can cause the upgrade to fail and can render the server unrecoverable.

To log off other remote users on Windows, use the following steps:

- a. Click Start → Settings → Control Panel → Administrative Tools → Terminal Services Manager.
- b. On the Users tab, right-click the users that should *not* be logged on to the server and select **Logoff** from the menu (Figure 5-15).

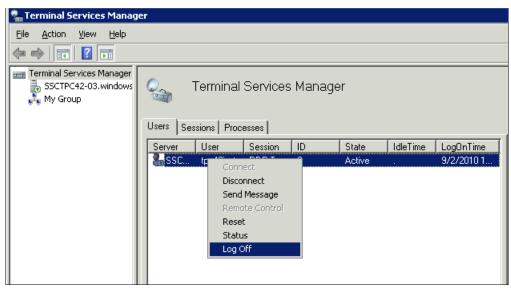


Figure 5-15 Terminal Services Manager

- c. Close the Terminal Services Manager window.
- 3. Stop all the Tivoli Storage Productivity Center services, Window, Linux, and AIX.

To stop the services on *Windows*, use the following steps:

- a. Go to Start \rightarrow Setting \rightarrow Control Panel \rightarrow Administrative Tools \rightarrow Services.
- b. Right-click the service and select **Stop**. The following services must be stopped:
 - IBM WebSphere Application Server V6 Device Server
 - IBM Tivoli Storage Productivity Center Data Server
 - IBM Tivoli Common Agent <directory> where <directory> is where the Common Agent is installed. The default is <TPC_install_dir>/ca.
 - IBM WebSphere Application Server v6.1 CSM if you also have Tivoli Storage Productivity Center for Replication

To stop the services on *Linux*, use the following steps:

Device server:

/<TPC_install_directory>/device/bin/linux/stopTPCF.sh

Data server:

/<TPC install directory>/data/server/tpcdsrv1 stop

- Common agent:

/<common agent install directory>/ca/endpoint.sh stop

Storage Resource agent:

/<SRA_install_directory>/agent/bin/agent.sh stop

- IBM WebSphere Application Server V6.1 - CSM:

/<usr or opt>/IBM/replication/eWAS/profiles/CSM/bin/stopServer.sh server1
-username <username> -password <passsword>

In this command, <username> represents the ID of the Tivoli Storage Productivity Center superuser, and represents the password for that user.

To stop the services on AIX, use the following steps:

Device server:

```
/<TPC install_directory>/device/bin/aix/stopTPCF.sh
```

- Data server:

```
stopsrc -s TSRMsrv1
```

Common agent:

```
/<common agent install directory>/ca/endpoint.sh stop
```

Storage Resource agent:

```
/<SRA_install_directory>/agent/bin/agent.sh stop
```

- IBM WebSphere Application Server V6.1 - CSM:

```
/<usr or opt>/IBM/replication/eWAS/profiles/CSM/bin/stopServer.sh server1
-username <username> -password <passsword>
```

In the command, <username> represents the ID of the Tivoli Storage Productivity Center superuser, and <password> represents the password for that user.

- 4. Back up your current Tivoli Storage Productivity Center V4.1 server and databases (TPCDB and IBMCDB). IBMCDB is the Agent Manager database and TPCDB is the Tivoli Storage Productivity Center database. This step is important in the event of an upgrade failure:
 - a. Back up your Tivoli Storage Productivity Center database using the DB2 backup process
 - b. For Tivoli Storage Productivity Center and Tivoli Integrated Portal single sign-on authentication configuration, back up the WebSphere configuration files. The configuration files are located in the following directories:

```
TIP_installation_directory/profiles/TIPProfile/bin
TPC_installation_directory/device/apps/was/profiles/deviceServer/bin
```

The backup file has the following name, (yyyy is year, mm is month, and dd is day): WebSphereConfig yyyy mm dd.zip

• Run the following commands on UNIX or Linux systems:

```
/IBM/Tivoli/tip/profiles/TIPProfile/bin/backupConfig.sh
/IBM/TPC/device/apps/was/profiles/deviceServer/bin/backupConfig.sh
```

• Run the following commands on Windows systems:

```
\IBM\Tivoli\tip\profiles\TIPProfile\bin\backupConfig.bat \IBM\TPC\device\apps\was\profiles\deviceServer\bin\backupConfig.bat
```

- c. Back up the following registries:
 - InstallShield registries:

```
Windows: C:\Program Files\Common Files\InstallShield\Universal\IBM-TPC
```

AIX: /usr/lib/objrepos/InstallShield/Universal/IBM-TPC/

UNIX: /root/InstallShield/Universal/IBM-TPC

• SRM legacy registry:

AIX (where # can be any number): subsystem TSRMsrv#

UNIX: /etc/Tivoli/TSRM

- Windows registry
- Common agent registry (if you have Data agents and Fabric agents installed):

Windows: C:\Program Files\Tivoli\ep*

AIX or UNIX: /usr/tivoli/ep*, /opt/tivoli/ep*

- d. Back up the Tivoli GUID setting. Go to the following locations:
 - Windows: C:\Program Files\Tivoli\guid
 - AIX or UNIX: /opt/tivoli/guid

In those locations, run the following command:

tivguid -show >tpc tivguid.txt

e. Back up the Agent Manager files and directories if you have Agent Manager installed:

```
AM_installation_directory/AppServer/agentmanager/config/cells/
AgentManagerCell/security.xml
AM_installation_directory/AppServer/agentmanager/installedApps/
AgentManager.ear/AgentManager.war/WEB-INF/classes/resources/
AgentManager.properties
AM_installation_directory/os.guid
AM_installation_directory/certs
```

f. Back up Tivoli Storage Productivity Center server files and directories.

```
TPC_installation_directory/config
TPC_installation_directory/data/config
TPC installation directory/device/config
```

g. Back up the Data agent and Fabric agent files and directories (if you have the Data agent and Fabric agent installed).

```
TPC_installation_directory/config
TPC_installation_directory/ca/cert
TPC_installation_directory/ca/config
TPC_installation_directory/ca/*.sys
TPC_installation_directory/ca/subagents/TPC/Data/config
TPC installation_directory/ca/subagents/TPC/Fabric/config
```

- Back up any interim fixes or work around code provided by Tivoli Storage Productivity Center support.
- 5. Restart all Tivoli Storage Productivity Center services.

To start the services on Windows, use the following steps:

- a. Go to Start \rightarrow Setting \rightarrow Control Panel \rightarrow Administrative Tools \rightarrow Services.
- b. Right-click the service and select **Start**. The following services must be restarted:
 - IBM WebSphere Application Server V6 Device Server
 - IBM Tivoli Storage Productivity Center Data Server
 - IBM Tivoli Common Agent <directory> where <directory> is where the Common Agent is installed. The default is <TPC_install_dir>/ca.
 - IBM WebSphere Application Server v6.1 CSM if you also have Tivoli Storage Productivity Center for Replication

To start the services on Linux, use the following commands:

- Device server:

/<TPC install directory>/device/bin/linux/startTPCF.sh

- Data server:

/<TPC_install_directory>/data/server/tpcdsrv1 start

Common agent:

/<common agent install directory>/ca/endpoint.sh start

Storage Resource agent:

/<SRA install directory>/agent/bin/agent.sh start

- IBM WebSphere Application Server V6.1 - CSM:

/<usr or opt>/IBM/replication/eWAS/profiles/CSM/bin/startServer.sh server1
-username <username> -password <passsword>

In this command, <username> represents the ID of the Tivoli Storage Productivity Center superuser, and <password> represents the password for that user.

To start the services on AIX, use the following command:

- Device server:

/<TPC install directory>/device/bin/aix/startTPCF.sh

Data server:

startsrc -s TSRMsrv1

- Common agent:

/<common agent install directory>/ca/endpoint.sh start

- Storage Resource agent:

/<SRA install directory>/agent/bin/agent.sh start

- IBM WebSphere Application Server V6.1 - CSM:

/<usr or opt>/IBM/replication/eWAS/profiles/CSM/bin/startServer.sh server1
-username <username> -password <passsword>

In the command, <username> represents the ID of the Tivoli Storage Productivity Center superuser, and <password> represents the password for that user.

Attention: If possible, reboot the Tivoli Storage Productivity Center servers. This action stops any remaining Tivoli Storage Productivity Center Java processes that might not stop in a timely manner. It is important for the Tivoli Storage Productivity Center Device server to stop and restart cleanly. If this action does not occur, a server reboot might be indicated.

6. Stop all Tivoli Storage Productivity Center jobs:

Stop all jobs including performance monitor jobs, system and fabric probe jobs, scan jobs, and other probe jobs.

Migrating of Tivoli Storage Productivity Center components

After the Tivoli Storage Productivity Center server and services are started, you can start with the migration of Tivoli Storage Productivity Center.

We use the same installation program used for installing the product. Depending on the components that are already installed on the system, various panels are displayed.

Important: On an AIX/Linux system, you must source the DB2 instance before starting an upgrade. See "Installing Tivoli Storage Productivity Center components" on page 142, for instructions on how to source the DB2 instance.

To perform the upgrade on a Windows computer, execute the following procedure:

- 1. Double-click the setup.exe file, which is located in the directory where you extracted the installation images.
- 2. Choose the language that must be used for installation and click **OK** (Figure 5-16).



Figure 5-16 Language selection panel

3. Read the terms in the license agreement and, if you agree, select I accept the terms of the license agreement, and then click Next to continue (Figure 5-17).

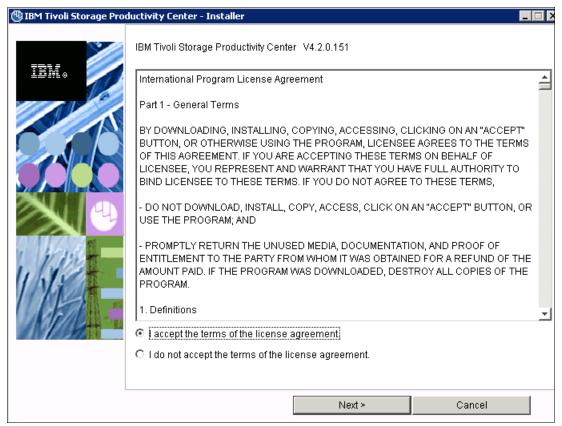


Figure 5-17 License agreement

- 4. In the next panel (Figure 5-18 on page 184), the following options are available:
 - Typical installation:

Use this option to upgrade all components on the same computer. The components are Servers, Clients, and Storage Resource Agent.

- Custom installation:

Use this option to select the components that you can upgrade.

Installation licenses:

Use this option to install the Tivoli Storage Productivity Center licenses. The Tivoli Storage Productivity Center license is on the DVD. You only need to run this option when you add a license to a Tivoli Storage Productivity Center package that has already been installed on your system.

The installation directory field is automatically filled with the Tivoli Storage Productivity
Center installation directory on the current system and disabled. In our case, a
previous version of Tivoli Storage Productivity Center is already installed in C:\Program
Files\IBM\TPC directory.

Select Custom Installation, and click Next to continue.

Tip: Select custom installation, which allows you to install each component of the Tivoli Storage Productivity Center separately.

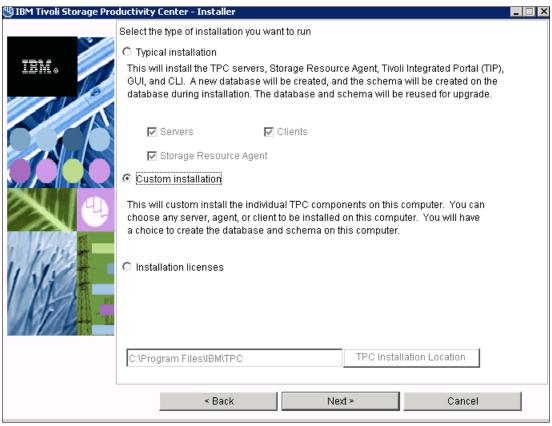


Figure 5-18 Tivoli Storage Productivity Center custom installation

5. The panel that lists the Tivoli Storage Productivity Center components opens (Figure 5-19). The components that have already installed on the system are discovered, selected for upgrade, and shown as disabled on this panel. The current version of each component is listed. In our case, we have a Tivoli Storage Productivity Center V4.1.1.55 installed on our system without local Data agents or Fabric agents. Click **Next** to proceed with the installation.

Tip: Storage Resource Agent will be upgraded using Tivoli Storage Productivity Center user interface after the Tivoli Storage Productivity Center upgrade.

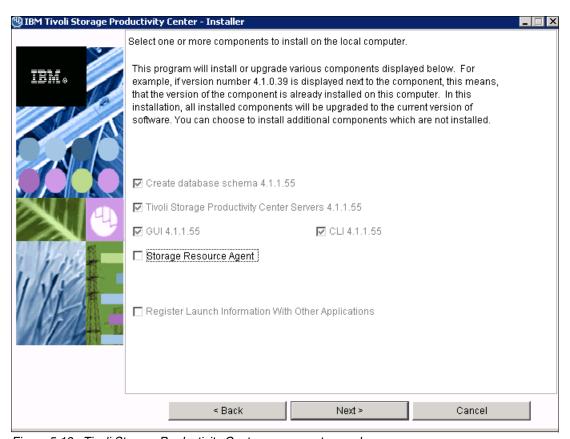


Figure 5-19 Tivoli Storage Productivity Center components panel

6. If you are running the upgrade on a system with at least 4 GB but less than 8 GB of RAM, a warning message is issued (Figure 5-20). Click **OK**.

Memory:

- ▶ 8 GB of RAM is the minimum memory requirement to run *both* Tivoli Storage Productivity Center and Tivoli and Tivoli Storage Productivity Center for Replication.
- ► If you have less than 8 GB of RAM, you have to run either Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication because of system load. To do that, you must disable Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication after installation.

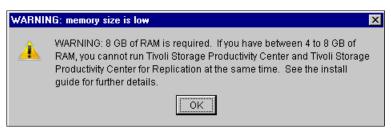


Figure 5-20 Tivoli Storage Productivity Center memory warning message

7. In the next panel (Figure 5-21), information in the DB2 user ID (administrator) and password fields is propagated. Click **Next** to proceed.

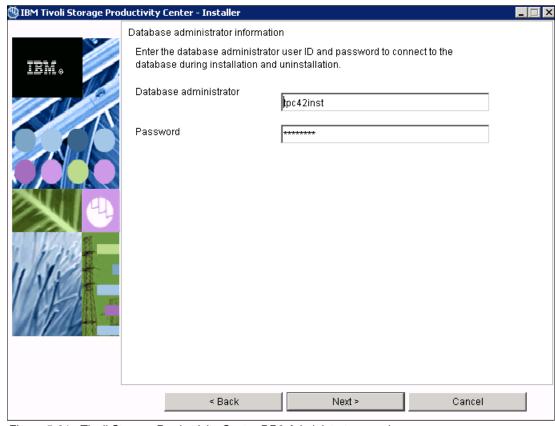


Figure 5-21 Tivoli Storage Productivity Center DB2 Administrator panel

8. The Database Schema panel is opens (Figure 5-22). All the information in this panel is already propagated. Verify it and click **Next** to continue.

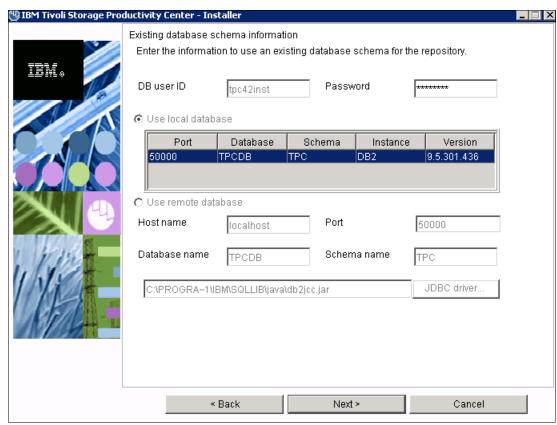


Figure 5-22 Tivoli Storage Productivity Center Database Schema pane

 In the Tivoli Storage Productivity Center servers panel (Figure 5-23), verify that the fields are filled with the correct information. The password fields are also filled with propagated information. Click Next.

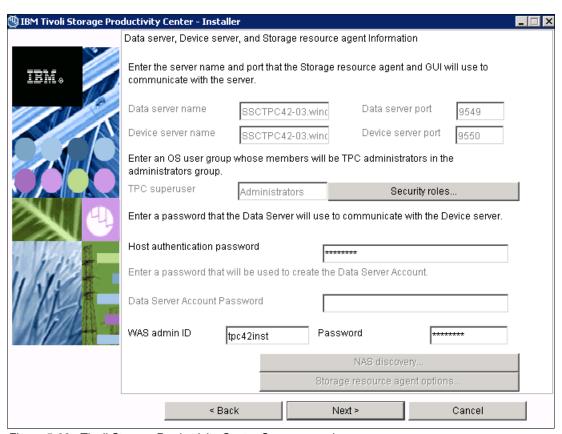


Figure 5-23 Tivoli Storage Productivity Center Servers panel

10.If Tivoli Storage Productivity Center detects that you have a DS8000, XIV, or SAN Volume Controller storage system, the Storage Subsystem Credential Migration Tool panel opens (Figure 5-24), which can help you to migrate the existing storage system credentials for the native interfaces. If you want to run the migration tool after the upgrade, clear the Run Storage Subsystem Credential Migration Tool check box. Otherwise, select the check box for this option. Click Next.

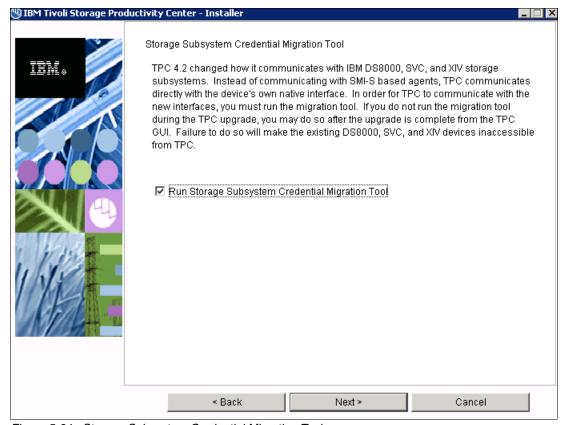


Figure 5-24 Storage Subsystem Credential Migration Tool

11.If the validation is successful, the summary panel opens (Figure 5-25). Review its content and click **Install** to start the upgrade.

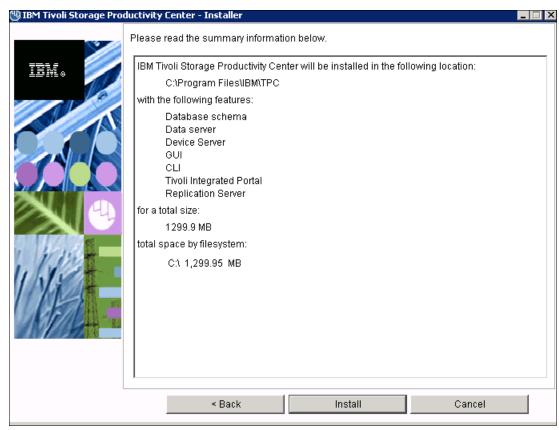


Figure 5-25 Summary panel

Tip: During the upgrade, you will not see the panels with Tivoli Integrated Portal installation.

The upgrade starts with deploying Storage Subsystem Credential Migration Tool (Figure 5-26).

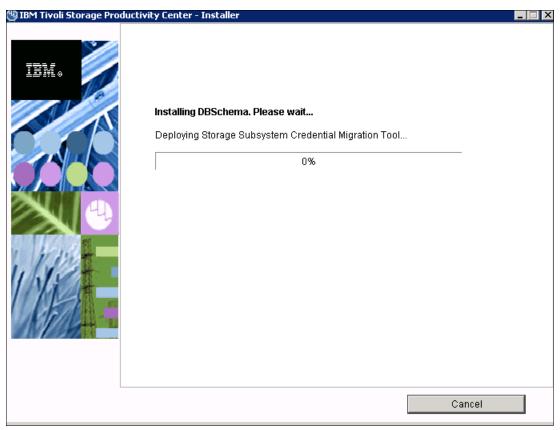


Figure 5-26 Deploying Storage Subsystem Credential Migration tool

12. The Storage Subsystem Credential Migration Tool panel opens (Figure 5-27). Use this panel to select the subsystems with credentials that can be updated automatically. Select the subsystems that you want to update, and then click **Update**.

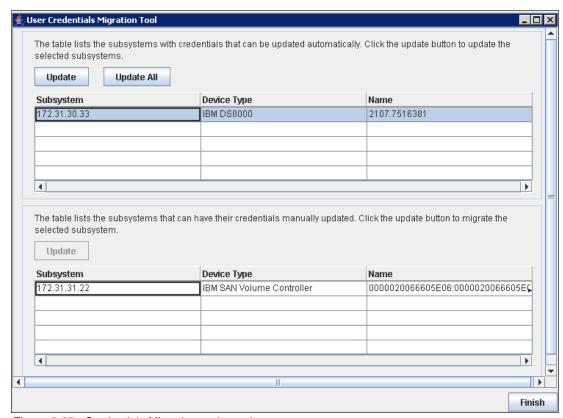


Figure 5-27 Credentials Migration tool panel

Updates: During this upgrade we update only the DS8000 subsystem. Updating SVC credentials is described in detail in "IBM SAN Volume Controller (SVC) or Storwize V7000" on page 230.

13. After you click Update, the subsystem is updated and removed from the table list. You can click Finish if you updated the selected subsystems, and click Yes to confirm and to close the Storage Subsystem Credential Migration Tool panel (Figure 5-28).



Figure 5-28 Confirmation to close Use Credentials Migration Tool panel

Multiple panels such as those shown in Figure 5-29 and Figure 5-30 are displayed and show the progress.

Tip: When you are upgrading the system, you might see several windows prompting you with the text **Replace Existing File**. Reply **Yes to All** to these prompts. Sometimes this dialog window is hidden behind the main installation panel. Be sure you look behind the main installation panel for any hidden dialog panels.

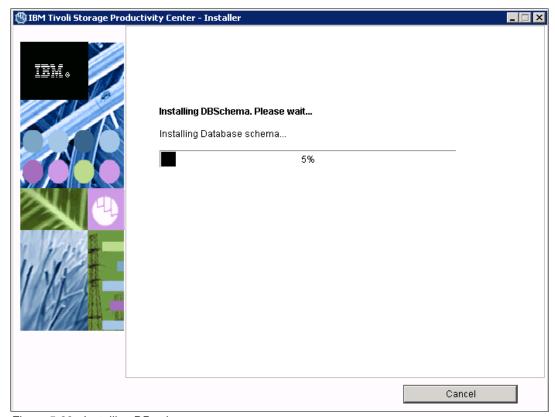


Figure 5-29 Installing DB schema

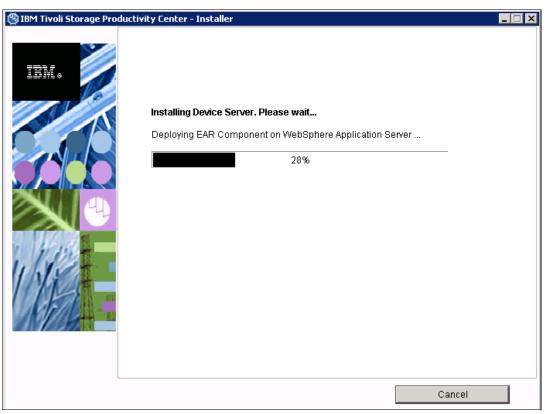


Figure 5-30 Installing Device Server panel

14. During the Tivoli Storage Productivity Center upgrade, the Tivoli Storage Productivity Center for Replication upgrade program is launched. The Tivoli Storage Productivity Center installation is temporarily suspended and remains in the background while the Tivoli Storage Productivity Center for Replication installation starts, and a welcome panel opens (Figure 5-31).

If Tivoli Storage Productivity Center for Replication is already installed on your system, it is upgraded. If it is not present, it is installed. In our system, we have a previous version of Tivoli Storage Productivity Center for Replication already installed, so the subsequent panels show a Tivoli Storage Productivity Center for Replication upgrade.

If it is the first time that Tivoli Storage Productivity Center for Replication is installed on the system, the installation process and panels are the same as those shown in the following chapters:

- Chapter 2, "Tivoli Storage Productivity Center installation on Windows" on page 21
- Chapter 3, "Tivoli Storage Productivity Center installation on Linux" on page 79
- Chapter 4, "Tivoli Storage Productivity Center installation on AIX" on page 125



Figure 5-31 Tivoli Storage Productivity Center for Replication Welcome panel

15. The installation wizard checks the system prerequisites to verify that the operating system is supported and the appropriate fix packs are installed. If the system passes the prerequisites check, the panel shown in Figure 5-32 is displayed. Click **Next** to continue.

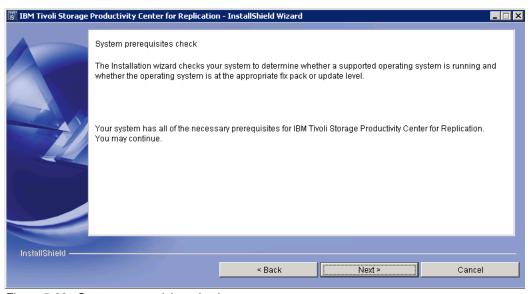


Figure 5-32 System prerequisites check

16. Read the terms in the license agreement panel (Figure 5-33), and if you agree, accept it and click **Next**.



Figure 5-33 License Agreement Panel

17. In the next panel (Figure 5-34), you can select the directory where Tivoli Storage Productivity Center for Replication will be installed. The directory where Tivoli Storage Productivity Center for Replication is correctly installed is proposed as the default location. You can accept it or change it based on your requirements. Click **Next** to continue.

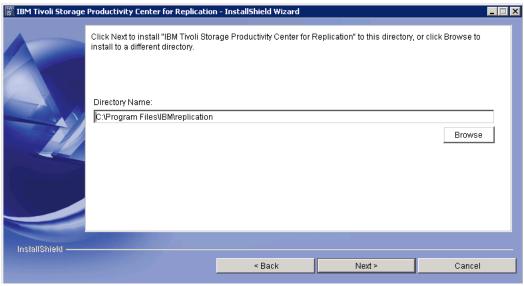


Figure 5-34 Tivoli Storage Productivity Center for Replication Installation directory

18. The upgrade program checks for currently running Tivoli Storage Productivity Center for Replication instances. If it is found, the message shown in Figure 5-35 is displayed.

Click **Yes** to continue the Tivoli Storage Productivity Center for Replication installation; Tivoli Storage Productivity Center for Replication service restarts during the upgrade.



Figure 5-35 Restart Tivoli Storage Productivity Center for Replication server during the upgrade

19. Review the summary of settings (Figure 5-36), and click **Install** to start the upgrade.

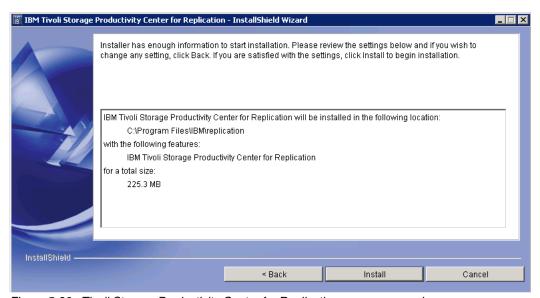


Figure 5-36 Tivoli Storage Productivity Center for Replication summary panel

The installation of Tivoli Storage Productivity Center for Replication starts. Several messages about the progress of the installation process are displayed, as shown in Figure 5-37, Figure 5-38, and Figure 5-39.

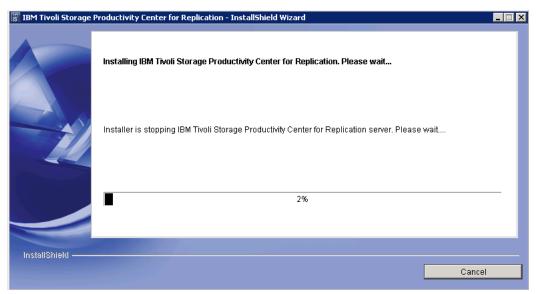


Figure 5-37 Stopping Tivoli Storage Productivity Center for Replication server

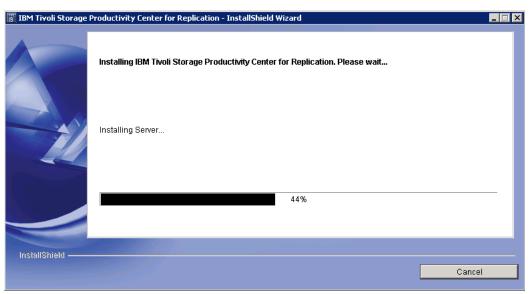


Figure 5-38 Installing Tivoli Storage Productivity Center for Replication

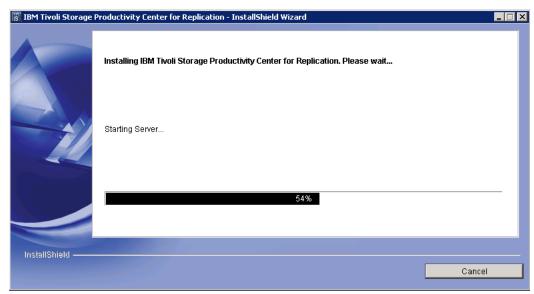


Figure 5-39 Starting Tivoli Storage Productivity Center for Replication

20. After the completion of the Tivoli Storage Productivity Center for Replication upgrade, the next panel opens (Figure 5-40). The panel also indicates the URL address to use for accessing the Tivoli Storage Productivity Center for Replication web interface. Click **Finish** to close the panel; the installation flows returns to the Tivoli Storage Productivity Center installation panels.

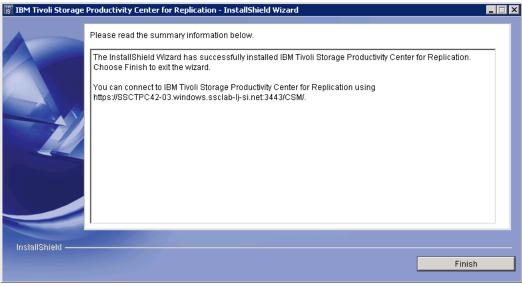


Figure 5-40 Tivoli Storage Productivity Center for Replication summary panel

21. The Tivoli Storage Productivity Center installation process continues, creating the uninstaller for Tivoli Storage Productivity Center and completes with summary information (Figure 5-41). Click **Finish** to complete the upgrade.

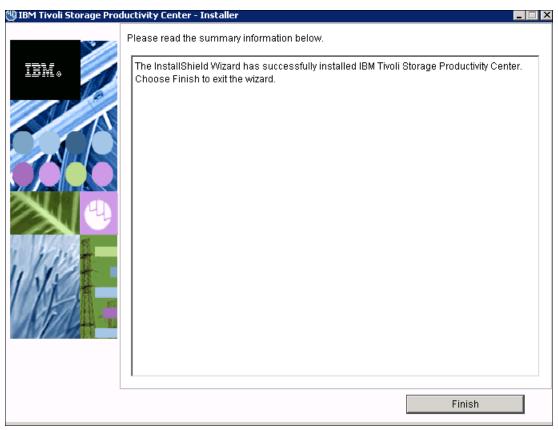


Figure 5-41 Tivoli Storage Productivity Center upgrade summary panel

Tip: If you are upgrading Tivoli Storage Productivity Center on AIX, you see the panel that indicates that the product is 100% installed, and you receive the following message, click **Yes to AII**:

/opt/IBM/TPC/service/service.sh exist on this system and is newer than the file being installed. Do you want to replace this file?

5.5 Upgrading Storage Resource Agent

You can upgrade Storage Resource Agent from Version 4.1 to Version 4.2 using one of the following methods:

- ► The Tivoli Storage Productivity Center installation wizard
- ► The Tivoli Storage Productivity Center user interface
- A Storage Resource Agent command line interface

When planning the upgrade of Storage Resource Agent, you must consider which agents can be migrated, which platforms and functions are unsupported, and what the limitations are. Details about Storage Resource Agent are described in Chapter 8, "Storage Resource Agent" on page 263.

In this section we show you how to upgrade Storage Resource Agent by using these methods.

Tip: Use the Tivoli Storage Productivity Center user interface to upgrade the Storage Resource agent because this method is the most common and typical. It provides you with more details about the installation and any failures.

If you use the Tivoli Storage Productivity Center installation wizard to do the upgrade, you also have to use the wizard if you want to uninstall the agent.

5.5.1 Installation wizard

When upgrading the Tivoli Storage Productivity Center server using the installation wizard, you can select to upgrade the Storage Resource agent. If you choose not to upgrade the agent as part of the server upgrade, you can launch the graphical installer at a later time to upgrade the agent.

Complete the following steps:

 After you start the Tivoli Storage Productivity Center installation wizard you can choose the Typical Installation or Custom Installation. In this section, we document Custom Installation. Click Custom installation (Figure 5-42). Click Next.

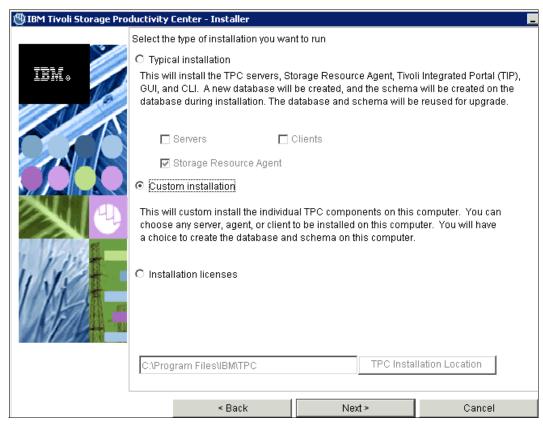


Figure 5-42 Tivoli Storage Productivity Center install wizard - Custom Installation

2. The panel shown in Figure 5-43 opens. Use this panel to select **Storage Resource Agent**. Click **Next**.

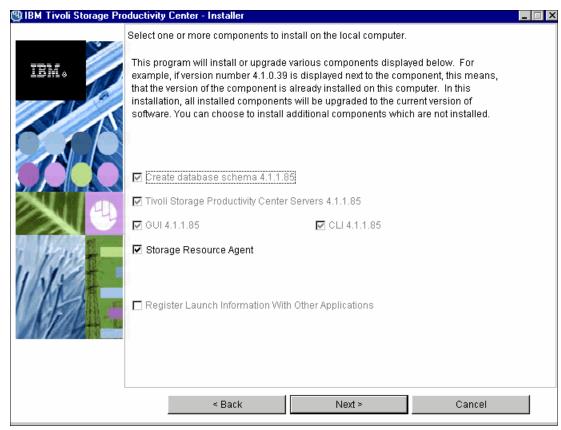


Figure 5-43 Installation wizard: selecting Storage Resource Agent

3. The Storage Resource Agent Information panel opens (Figure 5-44). You can enter the same options that are provided for Storage Resource Agent installation. For details, see Chapter 2, "Tivoli Storage Productivity Center installation on Windows" on page 21.

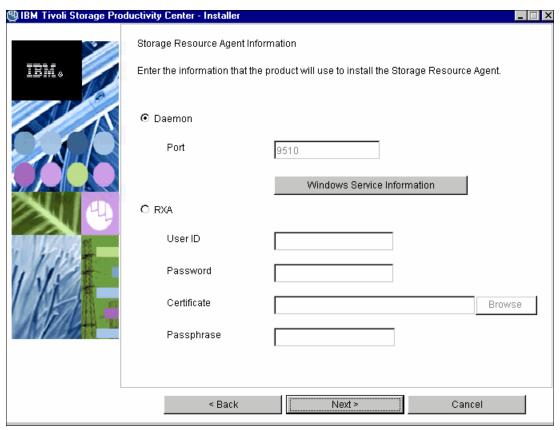


Figure 5-44 Storage Resource Agent information

With the successful upgrade of Tivoli Storage Productivity Center, the Storage Resource Agent is also successfully upgraded.

5.5.2 Tivoli Storage Productivity Center user interface

When upgrading Storage Resource agent using Tivoli Storage Productivity Center user interface, complete the following steps:

 In the navigation tree, click Administrative Services → Data Sources → Data/Storage Resource Agents (Figure 5-45).

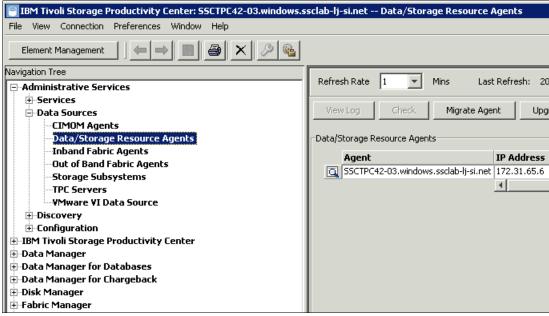


Figure 5-45 Selecting Data/Storage Resource Agent

2. In the content panel, select one or more agents for which you want to upgrade and click Upgrade Agents (Figure 5-46). The status of the Storage Resource agents that have to be upgraded are indicated by the message Need to upgrade agent software. If you have enabled the automatic upgrade action, the Storage Resource agent will be automatically upgraded after you upgrade Tivoli Storage Productivity Center server.



Figure 5-46 Upgrade Agents

3. The Create Storage Resource Agent Upgrade panel opens (Figure 5-47). Use this panel to select the computer and to schedule an upgrade of the Storage Resource Agent.

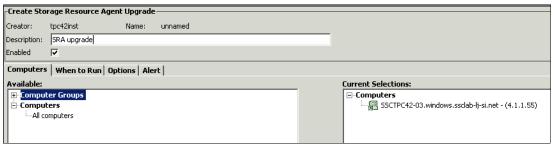


Figure 5-47 Create Storage Resource Agents upgrade

4. Run the upgrade job. The status of Storage Resource Agent is changed to Upgrading agent software (Figure 5-48).

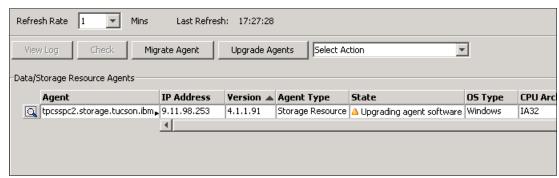


Figure 5-48 Upgrading Storage Resource Agent

5. Check the status in Job Management (Figure 5-49).

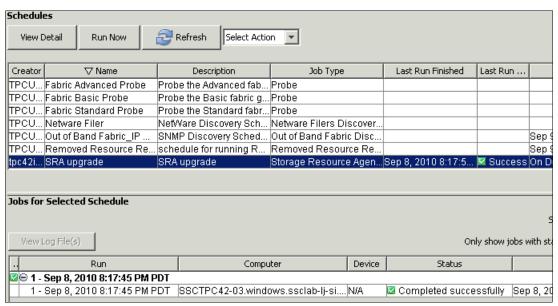


Figure 5-49 Job Management showing Storage Resource agent upgrade

After a successful upgrade, the Storage Resource Agent status is "Up" (Figure 5-50).

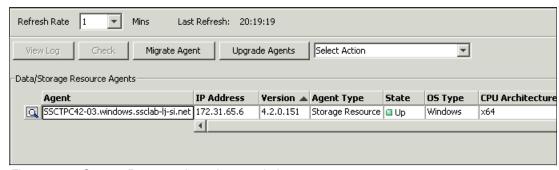


Figure 5-50 Storage Resource Agent is upgraded

Status: The Storage Resource agent upgrade job delivers the software upgrade packages to the agent computer. The job log displays the status of the delivery. The actual status of the upgrade is found in the agent log on the agent computer. If the agent log indicates that the upgrade failed and the state of the Storage Resource agent remains in the Upgrading agent software status, try restarting the agent and running the upgrade job again.

5.5.3 Command line interface

You can upgrade the Storage Resource agent manually by using the command-line interface. To upgrade the Storage Resource agent, complete the following steps:

1. Go to the location of the installation program (using the Storage Resource agent image) and go to the bin directory:

```
SRA image location/data/sra/operating system type/bin
```

2. From the bin directory, run the agent upgrade command:

```
Agent -upgrade -installLoc agent_install_directory
```

If the agent is run as a daemon service you must enter the following parameter and enclose the agent install directory name in quotation marks:

```
-commType Daemon
```

If the upgrade fails, you can check the return codes in the Tivoli Storage Productivity Center 4.2 Information Center by using the following link:

http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=/com.ibm.tp c_V42.doc/fqz0_r_return_codes_used_by_strg_resource_agent.html

5.6 Upgrading Tivoli Storage Productivity Center for Replication in a high availability environment

With Tivoli Storage Productivity Center for Replication V4.2 or later, an embedded database is used for the database repository. DB2 is no longer supported as the database repository for Tivoli Storage Productivity Center for Replication.

As a part of the upgrade process, an embedded repository is created and any data from an existing DB2 Tivoli Storage Productivity Center for Replication database is copied to the new database repository. This process is automatic and does not require any input or action.

If you are running Tivoli Storage Productivity Center for Replication in a high availability environment, you have to upgrade both active and standby Tivoli Storage Productivity Center for Replication servers.

If you are upgrading Tivoli Storage Productivity Center for Replication from DB2 to embedded repository, or you already have Tivoli Storage Productivity Center for Replication with embedded repository, the general procedures in a high availability environment are as follows:

- 1. Issue the takeover command to the standby server.
- 2. Upgrade the standby server.

Tip: Be aware that this action makes both the Tivoli Storage Productivity Center for Replication servers active.

- 3. Wait for the standby server to complete installation and then start it.
- 4. Upgrade the active server.

Attention: While upgrading the active server, avoid making any configuration changes to the sessions.

5. If no changes have been made to the configuration while the active server is being upgraded, issue a takeover command and reestablish the high availability function from the active server to the standby server. If configuration changes were made to the standby server, synchronize the high availability function from the standby server to the active server. Next, perform a takeover operation and reestablish the high availability function from the active server to the standby server.

During the initial synchronization, the current information in the database is saved and held until the synchronization is complete. If an error occurs during this process, the server database is restored to its original state before the synchronization process began. If an error occurs during the synchronization process and it causes the status to be in the disconnected or inconsistent state, you can reconnect to a synchronized state.

5.7 Upgrading System Storage Productivity Center (SSPC)

Tip: When you upgrade IBM Tivoli Storage Productivity Center Version 4.0 or later, IBM Tivoli Storage Productivity Center for Replication also is upgraded automatically.

To upgrade Tivoli Storage Productivity Center, follow these general steps.

Important: When you upgrade Tivoli Storage Productivity Center 4.1.1 to a release or patch, the DS Storage Manager code is removed from the following directory:

C:\Program Files\IBM\TPC\device\apps\was\profiles\deviceserver\InstalledApps\De faultNode\SMLiCJNLPGen.ear

After you upgrade Tivoli Storage Productivity Center, reinstall DS Storage Manager so that the launch-in-context feature of IBM System Storage DS4000 can function properly. To reinstall the application, use the DS Storage Manager CDs that were provided with SSPC.

- Perform the following steps to change the configuration of the Tivoli Integrated Portal so that it has the SSPC computer name that you want and not the SSPC computer name that was assigned by IBM Manufacturing.
 - Open a command prompt window and change the directory to:
 C:\Program Files\IBM\tivoli\tip\bin
 - Type wsadmin -lang jython -c AdminControl.stopServer('server1') and press Enter.
 - Type tipChangeHostName -h <hostname of machine> and press Enter.
 - Verify that the Tivoli Integrated Portal has stopped by opening C:\Program
 Files\IBM\tivoli\tip\profiles\TIPProfile\logs\server1\server1.pid. If Tivoli Integrated Portal
 is stopped, the file does not exist. If the file exists, it contains the process ID of Tivoli
 Integrated Portal which you can use in the Task Manager to stop Tivoli Integrated
 Portal.
 - Restart the SSPC server.
- 2. Start the Tivoli Storage Productivity Center installation program.
- 3. Select Typical installation.
- 4. Clear the check boxes for Agents and Register with the agent manager.
- The "User ID and password, and server and agent information" panel is displayed.
 Accept the defaults and click **Next**. If you changed the DB2 administrative password, you must enter the changed password here.
- 6. The "Summary Information" panel is displayed. Review the summary information panel and click Install.
- 7. The program will upgrade Tivoli Storage Productivity Center.



6

Agent migration and upgrade

In this chapter, we review the Common Agent Services and Storage Resource agent history. We describe various scenarios of how to upgrade the Common Agent Services to the Storage Resource agent. We also cover CIMOM to Native API (NAPI) migration guidelines.

6.1 Common Agent Services and Storage Resource agent history

Before introducing Storage Resource agents, Tivoli Storage Productivity Center used Tivoli Common Agent Services for software distribution and wanted state management. To take advantage of several Tivoli Storage Productivity Center software management features, the Common agent had to be installed on all managed endpoints. The Common agent environment consisted of the Common agent, Agent Manager, and Resource Manager.

The agents were used as programs that automatically performed some service, such as data collection. Tivoli Storage Productivity Center used Common Information Model (CIM) agents, Data agents, Fabric agents, and out-of-band Fabric agents to gather the data.

The Storage Resource agent was introduced in Tivoli Storage Productivity Center V4.1 as a lightweight agent to collect host disk and file system information. Now with Tivoli Storage Productivity Center V4.2 it includes full host monitoring functionality (including disk, file system, database, directory, and file information) for a broad set of platforms.

The Storage Resource agents do not require the Agent Manager and can be easily deployed to other systems using the Tivoli Storage Productivity Center GUI on the server system. The Storage Resource agents now perform the functions of the Data agents and Fabric agents.

Tivoli Storage Productivity Center V4.2 uses Storage Resource agents, CIM agents, and out-of-band Fabric agents to gather host, application, storage system, and SAN fabric information and send that information to the Data server or Device server.

Attention: Agent Manager can be used with Tivoli Storage Productivity Center V4.2 to allow communication with existing Data agents and Fabric agents that are present in the environment. However, no new functions were added to those agents for the V4.2 release.

The Tivoli Storage Productivity Center V4.2 installation program does not support installation of the Data agent or Fabric agent. If you want to install the legacy Data agent or Fabric agent, you must have a previous Tivoli Storage Productivity Center installation program that supports installing the Data agent or Fabric agent.

For optimal results when using Tivoli Storage Productivity Center, migrate the Data agents and Fabric agents to Storage Resource agents.

6.2 Benefits of Data agent and Fabric agent migrated to Storage Resource agent

Tivoli Storage Productivity Center provides the following benefits when you migrate an existing Data agent and Fabric agent to a Storage Resource agent:

► Fewer resources required. Storage Resource agents require fewer resources on a host computer than a Data agent or Fabric agent that is based on Common Agent Services. Additionally, you can deploy Storage Resource agent as a daemon or non-daemon service.

- ▶ Reduced complexity when deploying agents. You can deploy Storage Resource agents directly from the Tivoli Storage Productivity Center user interface and they do not require you to install Agent Manager. For Data agents and Fabric agents, you must use the Tivoli Storage Productivity Center installation program and ensure that Agent Manager is registered with the Data server and Device server.
- ► Improved interface when deploying, upgrading, and administering agents. You can manage Storage Resource agents using the nodes in the Administrative Services section of the Navigation Tree. You can deploy and upgrade Storage Resource agents on schedules that you define.

6.3 Prerequisites

When you plan the migration of Data agents and Fabric agents, consider which agents can be migrated, which platforms and functions are unsupported, and what the limitations are.

For the agent support, go to the following address:

http://www.ibm.com/support/docview.wss?uid=swg27019380#Agents

6.4 Scenarios to migrate from Common Agent Services to Storage Resource agent

In this section, we describe various scenarios of how to migrate the Data agents and Fabric agents to Storage Resource agents. The migration process installs a Storage Resource agent on a target host and then uninstalls the existing agent.

You can migrate Data agents and Fabric agents to Storage Resource agents by using one of the following methods:

- ► The Tivoli Storage Productivity Center installation wizard (for server system only)
- ► The Tivoli Storage Productivity Center user interface
- ► A command line interface

When you are upgrading Tivoli Storage Productivity Center using the installation wizard, the existing agents will be migrated as part of the Tivoli Storage Productivity Center upgrade. You can also use the wizard if you already have a Tivoli Storage Productivity Center V4.2 server installed.

When you have a Tivoli Storage Productivity Center V4.2 server installed, and are installing Storage Resource agent, consider the valid upgrade scenarios in Table 6-1.

Table 6-1 Agent upgrade scenarios on existing install

Tivoli Storage Productivity Center 4.2 installed	Use Tivoli Storage Productivity Center V4.2 installation program on server or to install a Storage Resource agent on a local computer	
Data agent or Fabric agent or both installed (V4.1.1 or earlier).	You can elect to migrate the Data agent or Fabric agent to a Storage Resource agent.	
Storage Resource agent is installed.	The Storage Resource agent is upgraded to 4.2.	
No agent is installed.	The default Storage Resource agent is installed.	

When you are upgrading a Tivoli Storage Productivity Center agent using the user interface, consider the valid upgrade scenarios in Table 6-2.

Table 6-2 Agent upgrade through the user interface

Tivoli Storage Productivity Center agent	Upgrade agent using Tivoli Storage Productivity Center 4.2 user interface
Data agent or Fabric agent or both on local computer	Not supported. You can migrate a Data agent or Fabric agent to a Storage Resource agent.
Storage Resource agent V4.1 on local computer	The Storage Resource agent is upgraded to the latest 4.2 level.
Storage Resource agent V4.2 on local computer	The Storage Resource agent is upgraded to the latest 4.2 level (must use force option).

When you are upgrading a Tivoli Storage Productivity Center agent using the command line interface, consider the valid upgrade scenarios in Table 6-3.

Table 6-3 Agent upgrade through the CLI

Tivoli Storage Productivity Center agent	Upgrade agent using Tivoli Storage Productivity Center 4.2 command line
Data agent or Fabric agent or both on local computer	Not supported. You can migrate the Data agent or Fabric agent to a Storage Resource agent.
Storage Resource agent 4.1 on local computer	The Storage Resource agent is upgraded to the latest 4.2 level (cannot change commtype).
Storage Resource agent V4.2 on local computer	The Storage Resource agent is upgraded to the latest 4.2 level (must use force option).

6.4.1 Installation wizard

You can migrate the Data agent or Fabric agent to the Storage Resource agent when you upgrade Tivoli Storage Productivity Center. To migrate the agents using the installation wizard, you only have to select the Storage Resource Agent; the upgrade procedure handles the upgrade of the agents.

Tips:

- ▶ When you upgrade the agents to a Storage Resource agent by using the local graphical installer, they must be uninstalled with the local graphical uninstaller.
- ► If you attempt to uninstall a Storage Resource agent that was installed using the local graphical installer from the Tivoli Storage Productivity Center GUI, the request will be denied.

6.4.2 Tivoli Storage Productivity Center user interface

You can migrate the Data agent or Fabric agent to a Storage Resource agent using the Tivoli Storage Productivity Center user interface. To schedule a migration job of the Data agent and Fabric agent to a Storage Resource Agent through the user interface, complete the following steps:

 In the Navigation Tree, click Administrative Services → Data Sources → Data/Storage Resource Agents. On the right, the agents and their state are listed. In our example, (Figure 6-1) the *state* shows that the agent must be migrated.

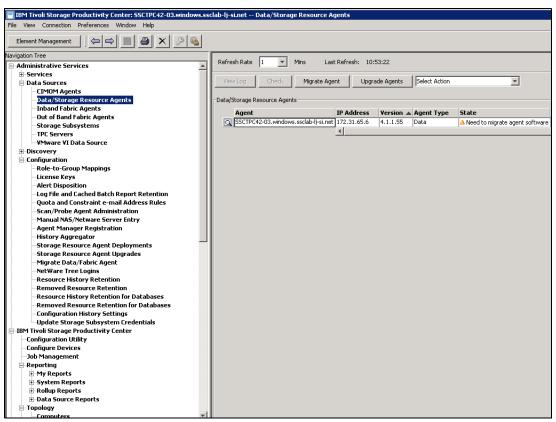


Figure 6-1 Selecting Data/Storage Resource agent

2. Select the agent that you want to migrate and click **Migrate**. The Create Data/Fabric Agent Migration panel opens (Figure 6-2).

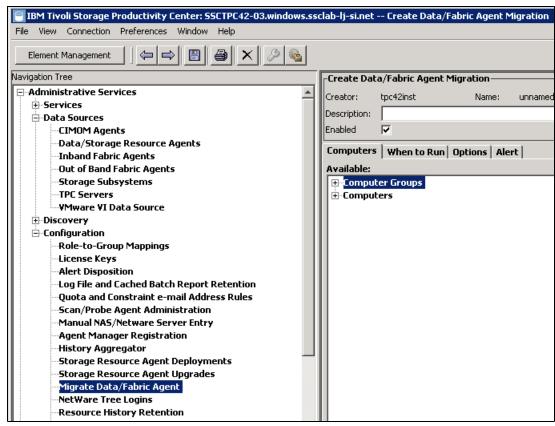


Figure 6-2 Migrate Data/Fabric Agent

3. The Computer selection tab (Figure 6-3) is where you select computers that have Data agents, Fabric agents, or both. Select the computer and schedule a migration job in the **When to Run** tab.

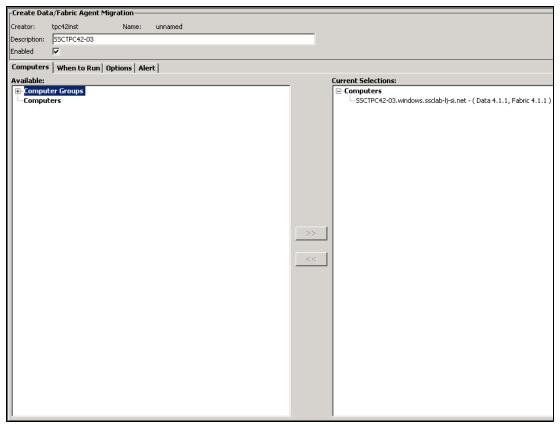


Figure 6-3 Computer selection

Tip: When a computer has both a Data and a Fabric agent, the migration job will always migrate both agents. There is no option to migrate one and not the other.

If both the Data agent and the Fabric agent are being migrated, the migration will migrate both to a single Storage Resource agent.

If the agent only has one agent, after migration the Storage Resource agent will be capable of performing both Data and Fabric functions. The concept of a well placed Fabric agent has been removed in this release.

4. In the Options tab, you can select how a Storage Resource agent is run after the migration. In our example we select to run the agent as a daemon service (Figure 6-4).

-Storage Resource Agent Runtime Operation				
✓ Use daemon service for runtime operation				
Run daemon service as specified user (Windows Only)				
User				
Password				
-Agent Migration	Options —			
Select the deployment method for agent migration				
Use comm	on agent services			
C Use SSH,	SMB/CIFS, RHOST, or RSH			
Allow the mi	gration to proceed even if the storage resource agent is unable to assume all the responsibilities cu	urrently held by the Data agent.		
-Credentials				
If the agent is not	operating as a daemon or the server is not using common agent services for migration, provide the	e credentials for use with the SSH,		
User				
Password				
Certificate Location	1			
Passphrase				

Figure 6-4 Storage Resource Agent Runtime operations

5. When you click **Save**, the panel shown in Figure 6-5 opens. The job is not saved until the verification is complete and you click **Proceed**.

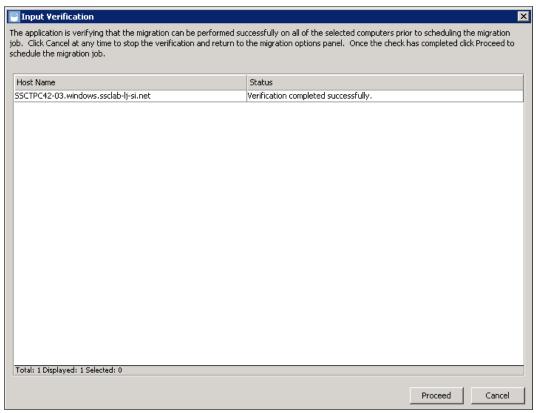


Figure 6-5 Agent input verification

The Save As dialog window opens (Figure 6-6).



Figure 6-6 Save migration job

6. Click **OK** and go to Job Management panel to check the status of migration job. In our scenario this is SSCTPC42-03. See Figure 6-7.

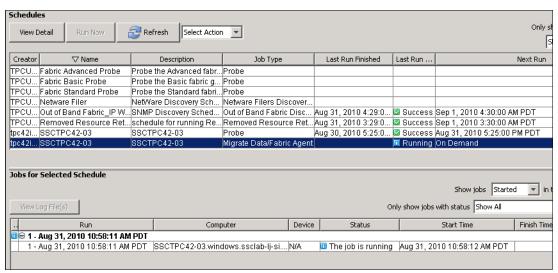


Figure 6-7 Migration job running

Tip: Each migration job will create one job log, regardless of how many computers are selected. When multiple computers are being migrated, the migrations are performed simultaneously in a maximum of 10 threads.

The progress of each computer can be tracked by host name.

If the migration completes with warnings, the migration succeeded although a minor issue exists.

7. In our example, the migration job completed with warnings (Figure 6-8) because the migration process was not able to cleanup some of the old files on the remote computer. This issue is common on Windows.

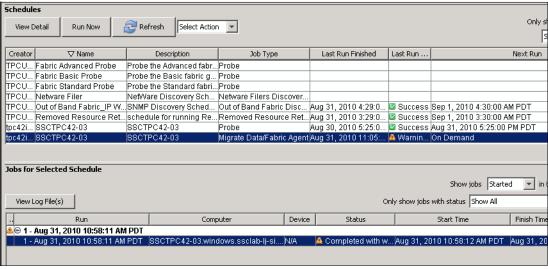


Figure 6-8 Migration job warnings

8. Click **View Log File(s)** to view the details (Figure 6-9).

```
8/31/10 10:58:12 AM STX03061: Job queued for processing. Waiting for idle thread.
8/31/10 10:58:12 AM STX03741: [SSCTPC42-03.windows.ssclab-lj-si.net] Deploying the migration image to SSCTPC42-03.windows.ssclab-lj-si.net.
8/31/10 10:58:19 AM SRX09721: [SSCTPC42-03.windows.ssclab-lj-si.net] The migration image has been deployed to SSCTPC42-03.windows.ssclab-lj-si.net.
8/31/10 10:58:19 AM SRX09721: [SSCTPC42-03.windows.ssclab-lj-si.net] Preparing the migration image on SSCTPC42-03.windows.ssclab-lj-si.net.
8/31/10 10:58:21 AM SRX098731: [SSCTPC42-03.windows.ssclab-lj-si.net] Starting the migration progress on SSCTPC42-03.windows.ssclab-lj-si.net.
8/31/10 10:58:21 AM SRX1002W: [SSCTPC42-03.windows.ssclab-lj-si.net] Starting the migration progress on SSCTPC42-03.windows.ssclab-lj-si.net] Starting the migration progress on SSCTPC42-03.windows.ssclab-lj-si.net] Starting the migration starting the starting the starting the starting the starting through through the starting through the starting through the starting through through
```

Figure 6-9 View Log File

 In our example, Common agent log files were not deleted. To finish the migration any files under the TPC_install_dir\TPC\ca directory can be manually deleted.

After the successful migration Storage Resource Agent status indicates "Up" (Figure 6-10).



Figure 6-10 Storage Resource Agent status

6.4.3 Command-line interface

You can migrate the Data agents and Fabric agents to Storage Resource agents by using the command-line interface. To migrate the Data agent or Fabric agent to the Storage Resource agent, complete the following steps:

- Go to the location of the installation program and go to the bin directory:
 SRA image location/data/sra/operating system type/bin
- 2. From the bin directory, run the migrate command, depending on daemon or non-daemon:
 - For daemon-based service, use the following command:
 Agent -migrate -commType Daemon -serverPort 9549 -debug max
 - For non-daemon service, use the following command:
 Agent -migrate -serverPort 9549 -userid myuserid -certFile mycertfile
 -passphrase mypassphrase -debug max

If the upgrade fails, you can check the return codes in the Tivoli Storage Productivity Center V4.2 information center at the following link:

http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=/com.ibm.tpc_V42.doc/fqz0_r_return_codes_used_by_strg_resource_agent.html

6.5 CIMOM to NAPI

When you upgrade to Tivoli Storage Productivity Center V4.2, a migration is required to switch to the Native API. The reason is that the XIV, SVC, and DS8000 Tivoli Storage Productivity Center V4.2 use only the Native API. The migration can be done prior to or during the install of Tivoli Storage Productivity Center V4.2, or even later, but you will not be able to use these devices until you complete the migration.

The three options for migrating the CIMOM user credentials/access information to NAPI are as follows:

- ▶ Before: You can provide authentication information while running earlier Tivoli Storage Productivity Center versions before upgrading to Tivoli Storage Productivity Center V4.2 by running the stand-alone credential migration tool. The information will be stored in the database for later use.
- ▶ During: During the upgrade, the installer determines whether you provided user authentication information for Native API devices. If you did not, the installer provides an option to launch the stand-alone credential migration tool.
- ► After: After upgrading to Tivoli Storage Productivity Center V4.2, you can select Administration Services → Data Source → Storage Subsystems and use the panel to provide new authentication information. The Configure Devices wizard will usually not work, because typically the Native API devices are already part of a probe job.

The credentials migration tool is described in Chapter 5, "Migrating Tivoli Storage Productivity Center base code to current level" on page 161.

Consider the following information:

- ► If you migrate a NAPI devices either prior to or as part of the upgrade to Tivoli Storage Productivity Center V4.2, any embedded DS8000 CIMOMs, SVC CIMOMs, and XIV CIMOMs will automatically be deleted from Tivoli Storage Productivity Center. Proxy DS CIMOMs will not be automatically deleted, even if Tivoli Storage Productivity Center knows of no other devices configured on that CIMOM.
- ► If the NAPI device is down at the time of the Tivoli Storage Productivity Center Data Server startup, its CIMOM will not be deleted.
- ► If you are upgrading from Tivoli Storage Productivity Center V4.1.1 to Tivoli Storage Productivity Center V4.2, and you want to migrate an existing Tivoli Storage Productivity Center V4.1.1 XIV CIMOM, previous historical data will be retained (true for all NAPI devices), but capability data will not be updated.

After the upgrade, a reprobe of the subsystem is necessary to enable new Tivoli Storage Productivity Center V4.2 capabilities (for example, creating and deleting XIV volumes).

Device configuration: Introducing the Native API

Tivoli Storage Productivity Center V4.2 provides a new access method to gather information from devices. This method is called the Native API and is at this time available for only a limited number of disk storage subsystems.

Although this chapter is focused on the Native API method, we also explain other new or changed parts of Tivoli Storage Productivity Center to provide a full picture of device configuration and handling within Tivoli Storage Productivity Center V4.2.

7.1 Native API and other changes

The Device Server has been reshaped in several ways. Although the most obvious change is the introduction of the Native API (NAPI), a full list of new and modified items is described in this section.

7.1.1 Full list of changes

New or modified items are as follows:

- XIV support has been enhanced in the following areas:
 - Added performance management
 - Added provisioning capability
 - Added alerts
- Enhanced discovery:
 - Supported subsystems discovered for NAPI even if they already are managed by a CIM Agent
 - Changes in the CIMOM discovery
- Credential migration tool to switch from using CIMOM to NAPI
- External processes for NAPI-based tasks controlled by the new External Process Manager
- ▶ Navigation Tree Changes (Figure 7-1 on page 223 shows an overview):
 - Added "Storage Subsystems" as a new entry under Administrative Services → Data Sources
 - Relabeled "Out of Band Fabric" option to "Switch and Subsystems (IP Scan)"
- ► New Configure Devices wizard (sometimes referred to as the *device configuration wizard*, although that is not the official name)

Figure 7-1 shows the Tivoli Storage Productivity Center Navigation Tree enhancements between V4.1 and V4.2

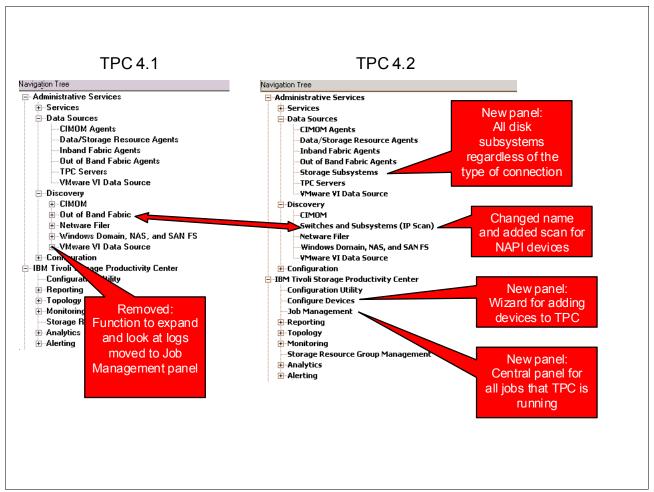


Figure 7-1 Navigation Tree changes between Tivoli Storage Productivity Center V4.1 and V4.2, overview

The remainder of this chapter has more information about the new or changed panels and tasks that are related to NAPI.

7.1.2 Changed panels and tasks

In terms of discovery, as well as adding and configuring devices in Tivoli Storage Productivity Center V4.2, be sure that you understand the summary in Table 7-1.

Table 7-1 Tasks and functions

Task	Function	
CIMOM discovery	The CIMOM discovery has been enhanced to filter out subsystems that are accessed through NAPI.	
NAPI discovery	This function has been combined with the existing IP scan under the Navigation Tree item named "Switch and Subsystem (IP Scan)".	
Storage subsystem panel	This panel is found under Administrative Services \rightarrow Data Sources \rightarrow Storage Subsystems and lists all disk storage subsystems, regardless of the method used for accessing the device. The advantage of this panel is that you can see which IP address is used to communicate with a device. In the table on the Disk Manager \rightarrow Storage Subsystems panel, you do not see that level of detail.	
Configure Devices wizard	The Configure Devices wizard is started when you click the add device button (for example Add Storage Subsystem on the Disk Manager → Storage Subsystems panel) from certain panels. You can also start the wizard by clicking the wrench icon in the icon bar:	
	This new wizard guides you through configuring devices with Tivoli Storage Productivity Center. Use this task when new devices have been discovered, or when you manually add devices. You can still do all the steps manually but the wizard can be more convenient by guiding you through the process.	

7.1.3 Behind the scenes: the External Process Manager

With the introduction of Native API, another architectural change has been introduced: the External Process Manager (EPM). This process manager is the link between the devices used by NAPI and Tivoli Storage Productivity Center. It is called External Process Manager, because now the jobs for the NAPI devices are started as external processes in the operating system, and are no longer running as threads within the Device server process. The advantage here is that the scalability and reliability is increased.

Figure 7-2 shows the high-level architecture of the EPM. You can see that the EPM starts external processes for each kind of device and each type of job.

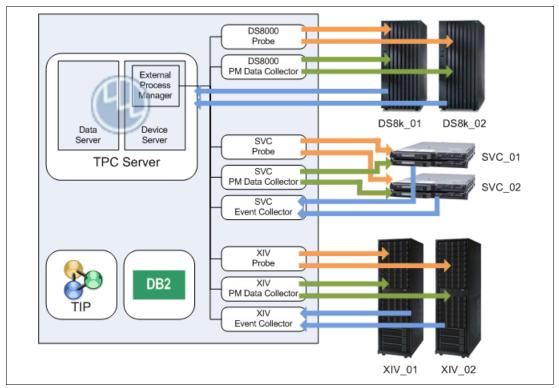


Figure 7-2 External Process Manager

Probes

With the implementation of EPM, additional changes exist in the way Tivoli Storage Productivity Center is doing probes for multiple devices. For every device type, there is a process running in the operating system. Each of those processes collects the information for one device at a time. As a result, the work runs in parallel for separate device types but runs sequentially for the devices of the same type. See Figure 7-3.

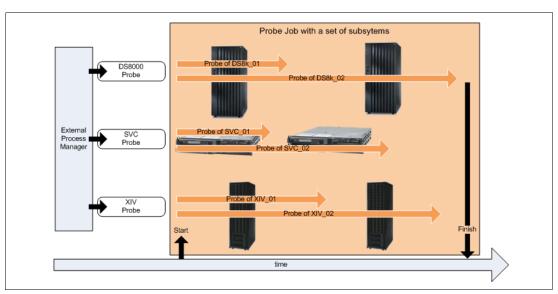


Figure 7-3 Running a Probe for multiple devices

Performance data

For the user, not much has changed with the introduction of the EPM. The minimum interval in which Tivoli Storage Productivity Center collects performance data is still five minutes for NAPI-attached devices. Thus you can expect to see, with each interval, one or more processes that are started to collect and insert the performance data into the Tivoli Storage Productivity Center database.

Although the ability to collect XIV performance data is new, that change is not caused by the introduction of the EPM.

In terms of stability, there has been a change in Tivoli Storage Productivity Center V4.2 which allows Tivoli Storage Productivity Center to fail over to a redundant path, such as a secondary Hardware Management Console (HMC), if it was not able to collect for some intervals. There is no parameter to control the retry and there will not be any alerts sent, but this greatly enhances the overall stability of performance data collections.

In Tivoli Storage Productivity Center V4.2.1, the failover mechanism has been added for CIM-based performance data collections.

Continuously running a performance data collection job

Because of the failover mechanism, letting a performance data collection job run continuously is now safer. We still believe that stopping and restarting it is a good approach, because it allows you to receive alerts if something fails. In a continuously running job, you do not receive any alerts if something fails.

Since the changes in Tivoli Storage Productivity Center V4.1.1, in which you can specify to run a job for 24 hours rather than 23 hours, there is little advantage in letting a job run continuously. By setting a job to 24 hours and restarting it daily, you do not lose any performance data and still have the chance to receive alerts about a failing job (at least once a day).

Figure 7-4 shows, in the red boxes on the left, that a duration length measured in hours allows for a daily restart. In the blue boxes on the right, a duration length measured in days will only let you restart it weekly (or monthly). Daily is not an available option here.

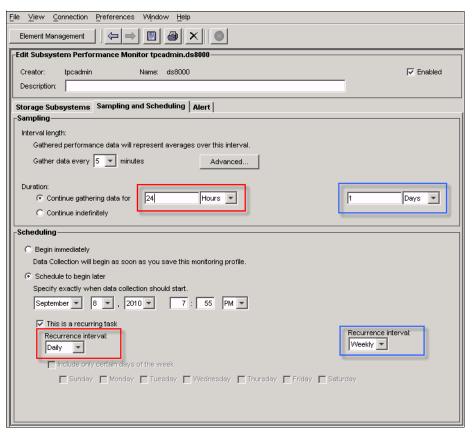


Figure 7-4 Differences between a duration of 24 hours and 1 day

7.2 Solution design for device access

This section is meant to help you with design planning. Read this section before you start implementing or configuring these functions. The outline follows this structure:

- When to use a function
- Considerations
- ▶ Requirements
- Implementation

7.2.1 Planning for NAPI and NAPI discovery

The Native API (NAPI) is a new way that Tivoli Storage Productivity Center uses to communicate with devices. The Native API does not replace CIM, SNMP or the in-band fabric interfaces. Although it is an addition to the ways Tivoli Storage Productivity Center can get information, you cannot decide which interface you want to use, because the support of NAPI is currently available for the following items:

- ► IBM System Storage DS8000 (DS8000)
- IBM System Storage SAN Volume Controller (SVC)
- ► IBM Storwize V7000 (Storwize V7000)
- IBM XIV Storage System (XIV)

Tip: The DS8000 support is limited to 2107 devices only, and does *not* include the family of like products such as DS6000 or Enterprise Storage Server.

The meaning of the name Native API immediately explains the difference to the other protocols. Native API uses the proprietary communication language of a device, and not a reduced set of standardized queries and commands.

When to use NAPI

For the devices we listed, Tivoli Storage Productivity Center V4.2 uses only the Native API. When you upgrade to Tivoli Storage Productivity Center V4.2, Tivoli Storage Productivity Center, an update/migration is required to switch to the NAPI, which can be done prior or during the installation. (It can be done later, but you will not be able to use such a device until you complete the migration). For that reason, the Supported Storage Products Matrix does not list any provider versions or interop namespaces for the IBM supported devices listed.

In addition to this new interface, the device server has been modified, so that together with the NAPI, the scalability and the reliability have been enhanced. Tivoli Storage Productivity Center is still not trying to replace the element management tools for those devices, but at the same time, customers have asked for better integration of IBM devices. As an example, for DS8000, specifying the logical subsystem (LSS) when provisioning volumes was not possible; this is now possible with Tivoli Storage Productivity Center V4.2.

The SMI standard will never include this level of detail because the intention of SMI-S is to abstract from the actual hardware devices.

Considerations for NAPI

The following list describes general considerations; we later describe more specific information for the supported devices:

- ► To discover NAPI devices automatically, you must use the Switch and Storage Subsystem (IP Scan) job, which was previously called Out of Band Discovery.
 - Because the discovery for NAPI devices does not return any status information about a device (as opposed to the CIMOM discovery), there is no real need to let the discovery run on a scheduled basis. However, the discovery of NAPI devices does not send any errors when a new device has been found, but credentials have not yet been provided to Tivoli Storage Productivity Center to connect to the device. As a result, leaving it enabled is not so annoying. With CIMOM discovery, a CIMOM discovered but without credentials provided yet can cause the CIMOM discovery job to fail.
- ▶ If new credentials are not available, monitoring of NAPI devices does not work after upgrading to Tivoli Storage Productivity Center V4.2. The CIM agent used prior to upgrading will not continue to be used.
- ▶ Probe performance is generally the same if not better than before.
- ► The discovery of new NAPI devices is part of the Switches and Subsystem (IP Scan) job. This job exists in earlier versions of Tivoli Storage Productivity Center but now has functions added so that it will identify subsystems that are used by the NAPI method. As long as there are not a lot of devices, we prefer to manually configure them or use the Configure Devices wizard rather than scan the IPs.

By default, Tivoli Storage Productivity Center has no subnets configured to be scanned. If you want to use it, be aware that you need to add the address range that you want to scan. If you do not want to continue to scan the IPs during future discovery jobs, you can remove that option after you have initially found the devices.

► The scan of IP ranges for subsystems and switches can be separated in such a way that Tivoli Storage Productivity Center either looks for switches, storage subsystems, both or neither. This setting is applied to all IP address ranges that you specify.

Figure 7-5 shows the Administrative Services \rightarrow Discovery \rightarrow Switch and Subsystem (IP Scan) panel where you can configure these settings. For example, in our environment, we specified the range 9.11.98.190 to 9.11.98.210.

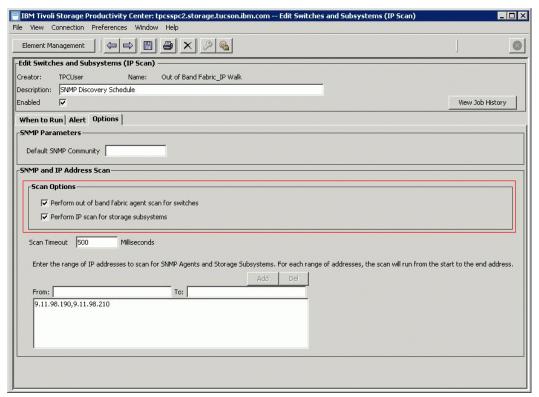


Figure 7-5 Choose for what kind of device Tivoli Storage Productivity Center will search

▶ If you want to be notified, change the alerts options for the IP Scan job, for example by entering your email address on the Alerts tab.

IBM DS8000

The DS8000 interacts with the NAPI as follows:

Access method used:

Enterprise Storage Server Network Interface (ESSNI)

► Failover:

For the communication with a DS8000 Tivoli Storage Productivity Center uses the ESSNI client. This library is basically the same library that is included in any DS8000 CLI. Because this component has built-in capabilities to do a failover from one HMC to another HMC, a good approach is to specify the secondary HMC IP address if your DS8000 has one.

The failover might still cause errors in a Tivoli Storage Productivity Center job, but the next command that is sent to the device should be using the redundant connection.

► Network:

No special network considerations exist. Tivoli Storage Productivity Center needs to be able to talk to the HMC as before, when the embedded CIMOM was used.

► Tivoli Storage Productivity Center is currently not able to provide specific messages for the vast majority of ESSNI error codes. You can still look up the errors in the DS8000 Information center; doing this often provides useful information (for example, that the user ID is wrong or that the password has expired), which will not be in any Tivoli Storage Productivity Center logs.

Consider the following example:

2010-08-05 16:58:09.296 HWNEP0003E A DS8000 ESSNI command failed. The error code is CMUN02021E

This message is the generic Tivoli Storage Productivity Center error; its action information directs the Tivoli Storage Productivity Center user to look up the ESSNI code in the DS8000 Information Center:

http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp

The information center reveals that this error code means "Unable to create logical volume: the volume number already exists."

IBM SAN Volume Controller (SVC) or Storwize V7000

The SVC and Storwize V7000 interact with the Native API as follows:

- ▶ The access method used is the Secure Shell (SSH).
- ▶ Failover:

For these devices, one node in the cluster is carrying out the role of the config node. This node manages the access through CIM, SSH and many other tasks, and it has a specific IP address. The cluster ensures that one node is always running as the config node. Therefore, from a Tivoli Storage Productivity Center perspective, the failover happens on the device itself.

Network:

Now that Tivoli Storage Productivity Center is accessing an SVC or Storwize V7000 cluster directly, you must be sure that Tivoli Storage Productivity Center can talk to the cluster. In earlier versions, when Tivoli Storage Productivity Center used a CIMOM, it only needed to talk to the CIMOM. If the CIMOM was running on the master console (SVC Version 4) it could use a separate physical network to communicate with the SVC cluster.

► Sessions:

The number of SSH open sessions that an SVC can have at one time is limited to 10 sessions. This limit of 10 sessions only includes external CLI access. The GUI and the embedded CIMOM in SVC Version 5 and later do not count against this limit.

You might also read something about 15 sessions. This amount is the number of new connections that can be opened per second. This number is bigger than the number of concurrent sessions because non-interactive (that is, script-driven) sessions can last less than a second per connection.

- Tivoli Storage Productivity Center handling of SSH keys for SVC or Storwize V7000 is as follows:
 - A default SSH key (tpc_svc.pem) is provided with Tivoli Storage Productivity Center.
 Although using only this key is convenient, it can compromise security, so it is best not to use it in an environment other than one used for testing or demonstrations.
 - Tivoli Storage Productivity Center can accept an SSH key in OpenSSH format or in PuTTY (.ppk) format. PuTTY keys will be automatically converted into OpenSSH format.
 - You can use passphrases only with OpenSSH keys.

- If the key is in PuTTY format and passphrases are not required, you must manually convert the key into the OpenSSH format.
- SVC and Storwize V7000 work with SSH keys as follows:
 - Public key is stored on the SVC or Storwize V7000 cluster.
 - · User or client application uses Private key.

If there are no keys uploaded that you want to use, you have three options:

- Use the default key that is included with Tivoli Storage Productivity Center (not desirable).
- Use a PuTTY key that you have generated and saved without a passphrase.
- Use an OpenSSH key that you have generated with or without a passphrase.
- ► Background information and general considerations for separate SVC versions:

Table 7-2 lists several general differences between SVC versions. Later in this chapter, we explain additional items to consider when you are adding an SVC to Tivoli Storage Productivity Center.

Table 7-2 Special considerations for SVC versions

Consideration	SVC 4	SVC and Storwize V7000 5+
Concept of Secure Shell (SSH) key	➤ SSH keys are associated to an authority level/role ➤ There are no individual users to associate a SSH key to Note: You can still upload multiple keys and let each user use a different key. This enables you to revoke access for a particular user, without any implications for other users.	 ► SSH keys are associated to a user ID. ► A user ID is always associated with a group within SVC or Storwize V70000 and therfore with an authority level. ► Because each key is associated with a user ID, you cannot use one keypair for more than one user ID.
Authority Level	The SSH key must have the Administrator access level.	 SVC Version 5 introduced a real user and group concept. The user must be part of the SVC Administrator group If the user has only monitoring access rights, you still can do probes, but you cannot run performance monitor jobs or perform any type of volume provisioning.
User ID/password	N/A	A user ID can be optionally assigned a password. This password is used only when a user wants to log in through the SVC or Storwize V7000 GUI or through the CIM interface; it is not used for SSH access.
SSH key upload with Tivoli Storage Productivity Center	 Performed through internal API Administrator user ID is required 	 Performed through CIM Administrator user ID is required

► As described in Table 7-2 on page 231, the association of SSH keys differs for SVC Version 4 and SVC or Storwize V7000 Version 5 and later. Figure 7-6 shows the logical difference.

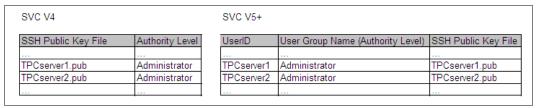


Figure 7-6 SSH key associations

Tip: Although SVC and Storwize V7000 Version 5 and later make use of user IDs, you must still start an SSH session to the SVC or Storwize V7000 with the user string **admin** but provide the key file of the user ID that you want to use for login. SVC and Storwize V7000 look through the list of key files to determine whether a matching public key can be found.

Using the **svcinfo 1suser** command, you can see which user ID is associated with the SSH session that you have open.

Unfortunately we could not find a command that lists all the stored keys and the corresponding user IDs.

► Guidelines:

- SSH key names: Give the SSH key files meaningful names, because determining which user is using a certain key pair is difficult to do later. For example, assign the user name as the file name to the key.
- Be sure that each Tivoli Storage Productivity Center server has its own pair of SSH keys when you work with an SVC or Storwize V7000. These keys can be used for accessing multiple SVCs or Storwize V7000s, but the association should always be as shown in Figure 7-7.

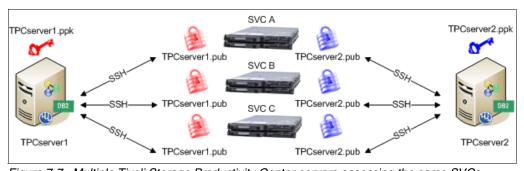


Figure 7-7 Multiple Tivoli Storage Productivity Center servers accessing the same SVCs

IBM XIV

The XIV interacts with the Native API as follows:

Access method used:

Native API is the XML-formatted version of the XIV command-line interface (XCLI).

Failover:

In Tivoli Storage Productivity Center V4.2.1, failover support is added for XIV devices. Tivoli Storage Productivity Center does not need all possible interface modules; instead Tivoli Storage Productivity Center queries the XIV during the setup for the IP addresses of the other interface modules.

Network:

Use only the address of one of the interface modules to add an XIV to Tivoli Storage Productivity Center, adding more than one IP address by starting the Configure Devices wizard again.

Requirements for NAPI

Firmware versions required are as follows:

- ▶ DS8000: Firmware Version 2.4.2 and later
- ► SVC: V4.2 and later (4.3 and later for Tivoli Storage Productivity Center for Replication)
- ► Storwize V7000: V6.1 and later (starting with Tivoli Storage Productivity Center 4.2.1)
- ► XIV: V10.1 and later

IP ports used from the native communication are as follows:

- DS8000: 1750SVC 4: 443
- ► SVC and Storwize V7000 5+: 5989 (to upload SSH keys) and 443 for normal operations
- ► XIV: 7778

Requirements for NAPI discovery

If you do not want to use the auto-discovery function to find your NAPI devices, you should manually add the devices or use the Configure Devices wizard (7.3, "Using the Configure Devices wizard" on page 241) to guide you.

To configure the auto-discovery, add the range to the IP addresses that Tivoli Storage Productivity Center should scan, as you did in previous versions of Tivoli Storage Productivity Center for out-of-band Fabric discoveries. You can configure this on the lower part of the Administrative Services \rightarrow Discovery \rightarrow Switch and Subsystem (IP Scan) panel (Figure 7-5 on page 229) where you add the IP address ranges.

Because there is no way to change the job name, we found it useful to change the job description, so that we were able to sort the list of jobs in the new Job Management panel.

Migration

Within the Tivoli Storage Productivity Center philosophy, the term *migration* is used when the architecture changes, for example going from CIMOM to NAPI. In contrast, Tivoli Storage Productivity Center refers to upgrades when only the version of a component changes, but the architecture stays the same, for example going from Storage Resource agent Version 4.1 to Storage Resource agent Version 4.2.

Migration methods and the credentials migration tool are described in "Credentials migration tool" on page 170.

Considerations

Consider the following information:

- ► If you migrate a NAPI device either prior to or as part of the upgrade to Tivoli Storage Productivity Center V4.2, any embedded DS8000 CIMOMs, SVC CIMOMs, and XIV CIMOMs will be automatically deleted from Tivoli Storage Productivity Center. Proxy DS CIMOMs will *not* be automatically deleted, even if Tivoli Storage Productivity Center knows of no other devices configured on that CIMOM.
- ► If the NAPI device is down at the time of the Tivoli Storage Productivity Center Data server startup, its CIMOM will not be deleted.
- ► If you are upgrading from Tivoli Storage Productivity Center V4.1.1 to Tivoli Storage Productivity Center V4.2, and you want to migrate an existing Tivoli Storage Productivity Center 4.1.1 XIV CIMOM, note the following information:
 - Previous historical data will be retained (true for all NAPI devices), but capacity data will not be updated.
 - After the upgrade, a reprobe of the subsystem is necessary to enable new V4.2 capabilities, such as creating and deleting XIV volumes.

7.2.2 Planning for CIMOM discovery

In this section we describe when to use CIMOM discovery, as well as considerations and requirements.

When to use CIMOM discovery

In most environments, we find that using CIMOM discovery does not have a large advantage, simply because most CIMOMs have security turned on, which means Tivoli Storage Productivity Center is unable to get a list of devices from the CIMOM.

As a result, the CIMOM discovery fails, leaving you a failed job log entry and potentially sending you an email about an obvious error.

Most of the time when you have a new device in your environment, you know about it. Because Tivoli Storage Productivity Center does not actually configure the device for you, we find little use of the automatic discovery of new CIM agents.

Considerations for CIMOM discovery

The CIMOM discovery of new CIMOMs has been changed in some ways to accommodate the implementation of NAPI. A CIMOM discovery can obtain limited information from a CIMOM even without authenticating, but in most cases this is not enough.

Here is a list of general CIMOM discovery considerations:

- ► CIMOM discovery is a process that serves three purposes:
 - Find new CIM agents.
 - Contact a known CIM agent to find new devices.
 - Get basic status information from devices managed by Tivoli Storage Productivity Center through CIM.
- ► Finding new CIM agents with the "Scan local subnet" option checked often results in a failed CIMOM discovery because it does not have credentials for those CIM agents found in the same subnet as the Tivoli Storage Productivity Center server. Add the credentials or remove the unwanted entry from the Administrative Services → Data Sources → CIMOM.

► On the Agents panel, disable the "Scan local subnet" option for future CIMOM discovery jobs, as shown in Figure 7-8.

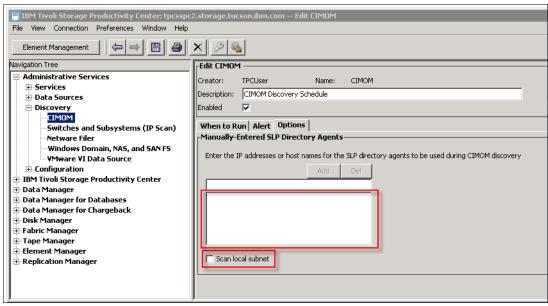


Figure 7-8 Disable scan of local subnet

➤ Tivoli Storage Productivity Center will discover CIMOMs that are not within the local subnet of the server by using Service Location Protocol (SLP), which must be configured at the Tivoli Storage Productivity Center side (provide the SLP Directory Agent (DA) IP address) and the SLP DA side (configure the list of devices available in that subnet).

Because the CIMOM discovery often fails for obvious reasons (for example, CIMOM that does not have credentials defined at the Tivoli Storage Productivity Center server), do not use the capability to look for new CIMOMs. To not use the capability, do not specify SLP DAs and do not let Tivoli Storage Productivity Center scan the local subnet looking for new CIMOMs, as shown in Figure 7-8. This way Tivoli Storage Productivity Center will still look for new devices at already configured CIMOMs and in addition you will still get status information.

When you read the following list, keep in mind that you might have a proxy CIMOM that has multiple devices of various types attached. However, Tivoli Storage Productivity Center might not support using all of those devices with CIM agents anymore because they are now accessed through the Native API.

- ► The discovery filters out CIMOMs of the devices that are supported only through the new NAPI interface, so the embedded CIMOMs for DS8000 and XIV are ignored, and so are any SVC CIM agents.
- ► If the discovery finds a DSopen API CIMOM, it is added to Administrative Services → Data Sources → CIMOM Agents list

The reason is that at this stage of a discovery, Tivoli Storage Productivity Center does not yet know what devices are attached to the newly discovered CIMOM. For it to know, you must add the credentials for the CIMOM, and then run the discovery again.

When Tivoli Storage Productivity Center can get a list of devices from the CIMOM, DS8000 devices will be filtered out, and the remaining (DS6000 and Enterprise Storage Server) devices will be added as managed devices.

If at this point, there are only DS8000 devices attached to the CIMOM, no managed devices will be added. The Tivoli Storage Productivity Center CIMOM will not be removed, because it can be used for other subsystems.

Requirements for CIMOM discovery

To use CIMOM discovery, there really are no requirements. The job is defined and activated by default, so there is little or nothing to do.

The modifications that you might want consider are as follows:

- Change the scheduling of the discovery.
- Add alert options to the job.
- ► Change the job description, for better sorting in the Job Management panel (see Chapter 13, "Job Management panel" on page 537 for additional details).

7.2.3 Planning for the Configure Devices Wizard

The Configure Devices wizard helps Tivoli Storage Productivity Center users add most types of data sources to Tivoli Storage Productivity Center by integrating the individual steps that you otherwise have to do manually. Data Sources you can add are as follows:

- Storage Subsystems
- ► Fabrics/Switches
- Computers
- Tape Libraries

Further simplification for configuring devices is provided by using Monitoring Groups (see 7.2.4, "Planning for Monitoring Groups" on page 238).

When to use the Configure Devices wizard

For experienced Tivoli Storage Productivity Center users, the Configure Devices wizard seems to be a small enhancement. After you start using it, you will realize that is also helpful for the experienced administrator.

The three scenarios for using the Configure Devices wizard are as follows:

- Scenario 1: Add and configure a data source.
 - When using the wizard in this case, it adds several steps to the beginning of the sequence (CIMON or NAPI discovery) that are not required in scenario 2.
- Scenario 2: Configure existing data sources.
 - In this case, you can configure a data source that has been discovered by Tivoli Storage Productivity Center but not yet set up.
- ► Scenario 3: Migrate existing data sources.
 - When you upgrade Tivoli Storage Productivity Center from a version earlier than V4.2, you must upgrade the connection information for NAPI devices. The various ways to do this step are described in Chapter 6, "Agent migration and upgrade" on page 209. However, all of those methods will eventually launch the Configure Devices wizard.

The major part of the Configure Devices wizard is the assignment of devices into monitoring groups. This ensures that the devices are put into probe jobs, and alerts are being defined.

If you do this step manually, the typical sequence of actions is as follows:

- 1. Run the scheduled discovery jobs.
- 2. If discovery finds a new data source, add credentials for that data source (such as a CIMOM).
- 3. Run discovery again to identify individual devices.
- 4. Check the discovery logs for any new found devices.
- 5. Define or add devices to a probe job.
- 6. Define alerts for the devices.

The wizard guides you through the whole process as these steps are performed in the background.

Considerations for the Configure Devices wizard

The following are considerations for the Configure Devices wizard.

- Because the wizard is a modal dialog box, you cannot switch back to the Tivoli Storage Productivity Center GUI until you finish or cancel out of it.
- ▶ By the time a device has been added to a probe job, it is considered *configured*, although a performance monitor job has not yet been set up. Remember, devices like tape libraries do not support performance monitoring.
- ► Certain data sources cannot be configured with the Configure Devices wizard. These include the following:
 - Data agents (replaced by Storage Resource agents)
 - Inband Fabric agents (replaced by Storage Resource agents)
 - Tivoli Storage Productivity Center Servers
 - VMware Hypervisor
 - Tivoli Storage Productivity Center agents running on an AIX VIO server
 - NetApp devices

Requirements for the Configure Devices wizard

There are no special requirements for using the wizard. It is available with any type of Tivoli Storage Productivity Center license.

Implementation

Table 7-3 shows a comparison of the ways for adding devices to Tivoli Storage Productivity Center. The purpose is to help you understand how the Configure Devices wizard guides you through the necessary steps for adding a device to Tivoli Storage Productivity Center.

Tivoli Storage Productivity Center V4.2 does not force you to use the new Configure Devices wizard (CD wizard) and the number of steps that you need to execute are not reduced. However, it will make sure that you run through this process quickly and in the right order.

In Table 7-3, the cells with bold borders indicate steps where you must supply the user credentials.

The cells that have a shaded background indicate the steps for which the Configure Devices (CD) wizard guides you, in contrast to the manual steps which are outlined in the first two columns.

Table 7-3 Comparison of the ways to add devices to Tivoli Storage Productivity Center

Not using the CD wizard for CIMOM devices	Not using the CD wizard for NAPI devices	V4.2 without discovery using CD wizard	V4.2 with CIMOM discovery using CD wizard	V4.2 with NAPI discovery using CD wizard
Create or ask for user credentials	Create or ask for user credentials	Create or ask for user credentials	Create or ask for user credentials	Create or ask for user credentials
			CIMOM discovery	
Add CIMOM to Tivoli Storage Productivity Center with credentials	Add NAPI to Tivoli Storage Productivity Center with credentials		Add credentials for CIMOM	
CIMOM discovery			Scheduled discovery	Scheduled discovery
		Add and configure device	Configure existing device	Configure existing device
		Add device with credentials		Add credentials
		Wizard runs discovery		Wizard runs discovery
		Select discovered devices	Select discovered devices	Select discovered devices
Add or define probe	Add or define probe	Add or define probe	Add or define probe	Add or define probe
Define alerts	Define alerts	job; define alerts	job; define alerts	job; define alerts

7.2.4 Planning for Monitoring Groups

Using Monitoring Groups is not a new concept in Tivoli Storage Productivity Center V4.2. You could define a Monitoring Group in previous versions of Tivoli Storage Productivity Center, but there was little documentation and few use cases that described the underlying concept and why you would want to use them.

A Monitoring Group is simply a collection of devices and has a distinct name. For every kind of monitoring operation, you can use those groups instead of adding each individual device to a job or alert.

In Tivoli Storage Productivity Center V4.2, this concept has been applied and three categories have been defined to group the devices together, so that you have fewer jobs and do not need to set up many alerts manually. For the various device types, such as Computers, Storage Subsystems, Fabrics, and Tape Libraries, the following groups are defined by default:

- ▶ Basic
- ► Standard
- Advanced

The context of Monitoring Groups, Jobs, and Alerts is shown in Figure 7-9.

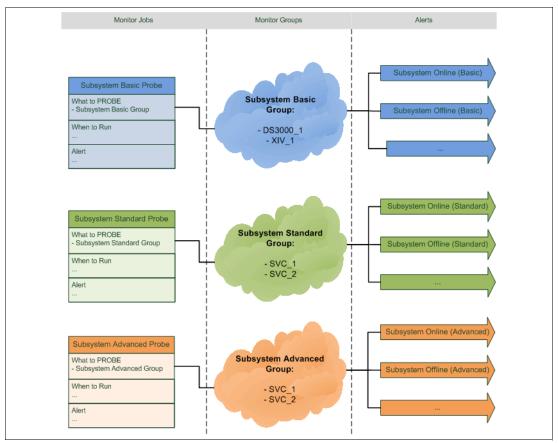


Figure 7-9 Context of Monitoring Groups, jobs, and alerts

When we use the term *groups*, we refer to the complete concept as shown in the figure. If we use the term *Monitoring Group*, we refer only to the real groups and not the concept with probes and alerts.

When to use Monitoring Groups

You do not have to decide whether to use Monitoring Groups, because they are predefined in Tivoli Storage Productivity Center. Although you can ignore the groups and create your own jobs as in previous releases, the predefined objects are shown on the Tivoli Storage Productivity Center GUI.

Considerations for Monitoring Groups

You should understand the following considerations to implement or customize Monitoring Groups:

- You can change the groups, jobs, and alerts that are provided by default.
- You cannot delete or rename the groups, jobs, and alerts that are provided by default.
- ▶ Although you can create additional groups and use them for additional jobs and alerts, or also associate them with existing jobs, this approach is not desirable because it can make understanding the relationship of the groups, jobs, and alerts more difficult.
- ► Although several alerts are defined, no one will receive these alerts. If an alert is triggered it is stored in the IBM Tivoli Storage Productivity Center → Alerting → Alert Log.

- ► The Monitoring Groups are predefined only for probes and alerts, but you can also use the groups for other jobs, especially in Data Manager.
- Monitoring Groups cannot be used for Subsystem Performance Monitor jobs.
- ► The most important difference between the basic, standard, and advanced groups is the probe frequency. The actual frequency depends on the device type, so for example not all advanced groups have the same frequency. Table 7-4 provides an overview of the initial frequencies, but you can change them according to your needs.

Table 7-4 Probe frequencies

Entity	Basic	Standard	Advanced
Computer	Monday, Wednesday	Daily	Every 12 hours
Storage Subsystem	Tuesday	Monday, Wednesday	Daily
Fabric	Monday, Wednesday	Daily	Every 12 hours
Tape Library	Sunday	Sunday, Wednesday	Daily

IBM Tivoli Storage Productivity Center Installation and Configuration Guide, SC27-2337 has details of the default settings for the probe jobs and alerts in the section "Alerts and schedules associated with monitoring groups and templates."

Requirements and implementation of Monitoring Groups

There are no real requirements for using the Monitoring Groups concept, but to fully apply the concept, you should implement the following changes:

- ► Change the alert options for the jobs and alerts, for example add your email address to receive email in case a jobs fails or an alert is triggered.
- ► Change the timing of the probe jobs.

You can find the definitions of the default Monitoring Groups in various places depending on which component of Tivoli Storage Productivity Center you are using:

- ▶ Data Manager → Monitoring → Groups → Computer
- ▶ Disk Manager → Monitoring → Groups → Subsystem
- ► Fabric Manager → Monitoring → Groups → Fabric
- ► Tape Manager → Monitoring → Groups → Tape Library

7.2.5 Planning for events

The approach of Tivoli Storage Productivity Center is not to replace any element manager software or SNMP manager, but rather to abstract the devices to a higher level. In addition to the use of the Native API, the SMI specification serves this purpose, both for getting data from a device and managing a device. It also introduces the concept of CIM indications for health and fault management.

CIM indications

CIM indications are a way to let a CIM client application, such as Tivoli Storage Productivity Center, know that something has happened on the device. The client can subscribe to receive indications for the events by supplying an address to which the indications will be sent (indication listener). This is one reason why Tivoli Storage Productivity Center should be installed only on servers with one NIC, or be configured as described in Chapter 3 of the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide*, SC27-2337, in the section titled "Configuring IBM Tivoli Storage Productivity Center with multiple IP addresses."

In addition to CIM indications, a CIM client can also poll the device for its health and operational status. The underlying idea of those two properties is that while one power supply of a device might fail (note that the health status would be showing this) it still can operate without any impact. All values are online and the performance is not affected.

Because SMI-S is meant to be a general specification, not device-specific, the health and operational status is translated and categorized by a CIM agent from the device internal event into a generic value. The advantage of this is that it enables a CIM client, such as Tivoli Storage Productivity Center, to display a status for a device without the need to load some kind of device description files, such as loading MIB files for an SNMP management application.

Events and the Native API

With the introduction of Native API and the External Process Manager (EPM), several event changes are introduced:

- ▶ The concept of CIM indications does not exist for the Native API.
- ▶ DS8000 asynchronously sends events to Tivoli Storage Productivity Center.
- ► For SVC, Storwize V7000, and XIV, Tivoli Storage Productivity Center will poll every minute to get new events from the subsystems. Therefore, each minute, you see processes in your operating system being started and stopped.

The goal of Tivoli Storage Productivity Center V4.2 is to retain the same level of device support of previous levels of Tivoli Storage Productivity Center. For this reason, no new alerts have been added.

7.3 Using the Configure Devices wizard

In this section, we provide a general overview of the steps the wizard uses to guide you through the process. We show you the panels to add a device. Most often, we use the term *device*, but the information applies to subsystems, fabrics and switches, computers, or tape libraries.

Generally, the wizard can be used to add a new device to Tivoli Storage Productivity Center or set up an existing device. For both processes, only the first part of adding a new device or selecting an existing device differs. After this step has been done, the remainder of the wizard guides your through the same steps of selecting a monitoring group for the probe and alerts.

From a licensing perspective, the Configure Devices wizard is aimed at the Basic Edition, so every Tivoli Storage Productivity Center customer can use the wizard. At the same time, the wizard does not set up performance monitors for you (performance monitors are not supported in the Basic Edition).

All this is similar for all types of devices, such as storage subsystems, fabrics or switches, computer systems (meaning computers with an installed Storage Resource agent), and tape libraries.

Considerations:

- ► Tivoli Storage Productivity Center considers any device that is not part of a probe job as unmonitored, and Tivoli Storage Productivity Center will list it in the existing devices that can be configured.
- ▶ Because every computer with an agent installed is automatically part of the default monitoring group, there are only a few situations in which the Configure Devices wizard presents a list with computers that are not monitored.
- For example, computers that are not monitored can occur if you moved a computer into another monitoring group, but that group is not associated with a probe job.

7.3.1 Scenario

Here we present the high-level steps of the Configure Devices wizard. We choose to not include the details of how you add a new device to Tivoli Storage Productivity Center at this point because that would make the list long and more complex to understand. On the left side of the windows, the wizard shows you the completed, the current, and the remaining steps.

We do not provide figures for every panel, for example, certain panels are simply confirmation panels that open only in certain circumstances. We do provide figures of panels that require you to actively do an action.

You can start the Configure Devices wizard from various locations within Tivoli Storage Productivity Center:

- ► Wrench Icon on the icon bar
- ► Administrative Services → Data Sources → Storage Subsystems → Add
- ▶ Disk Manager → Storage Subsystems → Add
- ► Fabric Manager → Fabrics → Add
- ► Tape Manager → Tape Libraries → Add
- Select device type

The windows we describe are as follows:

1. Select Device Type:

After you start the wizard, the Select Device Type panel opens (Figure 7-10). In this panel, select the device to be added.

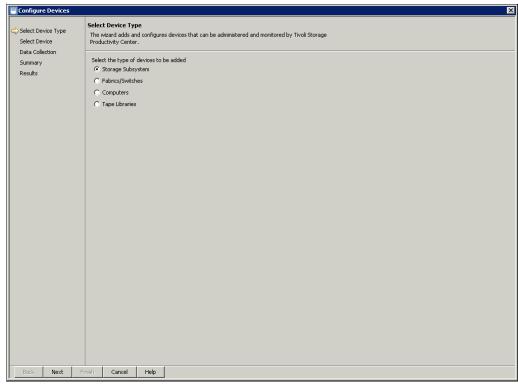


Figure 7-10 Select device type

2. Select device:

During this step you either add a new device or select an existing device that has been added to Tivoli Storage Productivity Center but has not yet been added to a probe. If the device has just been discovered, you might also have to add the credentials for that device. Figure 7-11 shows an example of selecting multiple devices.

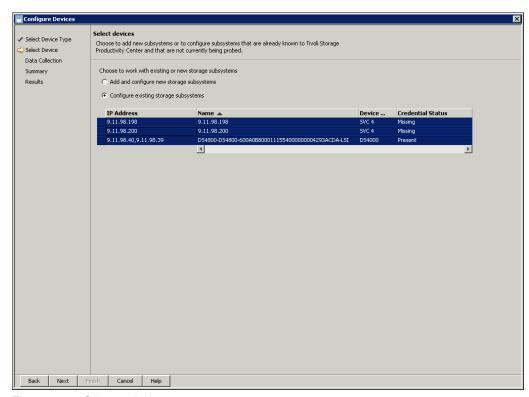


Figure 7-11 Select a device

If you select one or more devices, all the devices that you configure in a single step through the wizard will share the same probe schedule and alert conditions.

If at this point you choose to add a new device to Tivoli Storage Productivity Center (in the example, it would be a storage subsystem) on the next window you see a panel where you can select the device type and enter the required information for that device. We dedicate separate sections for this procedure later in this chapter because we want to simply provide a high level overview in this section. For more information, see the following sections:

- DS8000: 7.3.2, "Adding or configuring an IBM DS8000" on page 251
- SVC: 7.3.3, "Adding or configuring an IBM SAN Volume Controller (SVC) or Storwize V7000" on page 252
- Storwize V7000: Same as SVC See 7.3.3, "Adding or configuring an IBM SAN Volume Controller (SVC) or Storwize V7000" on page 252
- XIV: 7.3.4, "Adding or configuring an IBM XIV" on page 254
- CIMOM: 7.3.5, "Adding or configuring a CIMOM" on page 255
- Fabric/switches: 7.4, "Adding or configuring fabrics and switches" on page 255

After this step is finished, Tivoli Storage Productivity Center will initiate a short discovery job to verify it can talk to the device. After the discovery has completed, you can click finish to continue with the next step.

3. Configure devices:

Because at this step we have selected two SVCs which have been discovered by the Switch and Subsystem (IP Scan) job but no credentials have been provided, the wizard adds this step to the sequence.

You must update the credentials for the devices or remove them from the selection. In the example in Figure 7-12, we have updated the credentials for the SVC with the IP address 9.11.98.198 and, as a result, the credential status has changed from Missing to Present. The next step is to either do the same for the other SVC with the IP address 9.11.98.200 or remove that SVC from the list of selected devices.

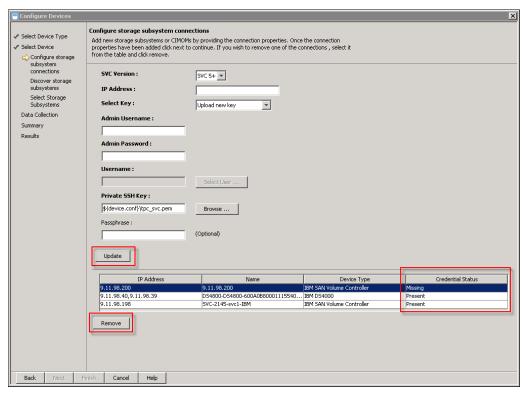


Figure 7-12 Configure devices

For this scenario, we have removed the second SVC. A confirmation page opens (Figure 7-13). At this step, you can still remove a device from the list of selected devices.

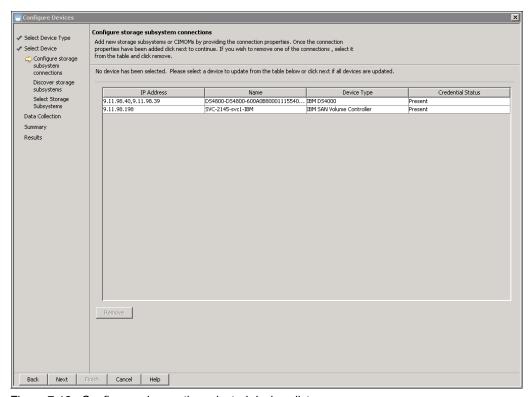


Figure 7-13 Confirm or change the selected devices list

If you click **Next**, Tivoli Storage Productivity Center initiates a discovery to verify the connection to the device, or in case you added a CIMOM, it looks for devices at that CIMOM.

If you click **Cancel**, you must confirm the cancellation (Figure 7-14).

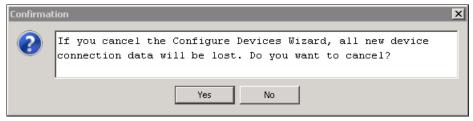


Figure 7-14 Cancel before discovery

After the discovery is finished, canceling the wizard results in a another confirmation message (Figure 7-15).

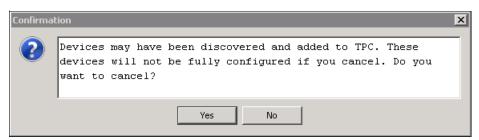


Figure 7-15 Cancel after successful discovery

4. Data collection:

On the panel in Figure 7-16, you must choose a monitoring group or template. This step adds your device to a group; because the groups are part of the predefined probe jobs and alert definitions, your device will be managed with those settings.

If you select **Monitoring Group** in the first drop-down list, the second drop-down menu lists all the defined groups that are associated with a probe job.

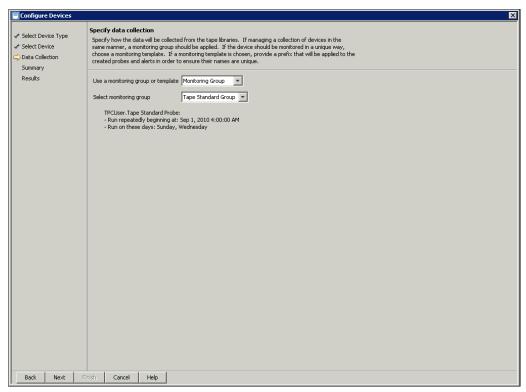


Figure 7-16 Data collection with Monitoring Groups

Tip: Although the group you select adds your device to a probe job and also to the alerts of the group, the alerts are not shown on this page. The wizard shows the alerts on the summary panel.

If you select **Monitoring Template**, as shown in Figure 7-17, you can create probe jobs and alerts. The wizard does not need all the detailed inputs, which you usually need to create new job. It derives that information for the three templates that are provided with Tivoli Storage Productivity Center.

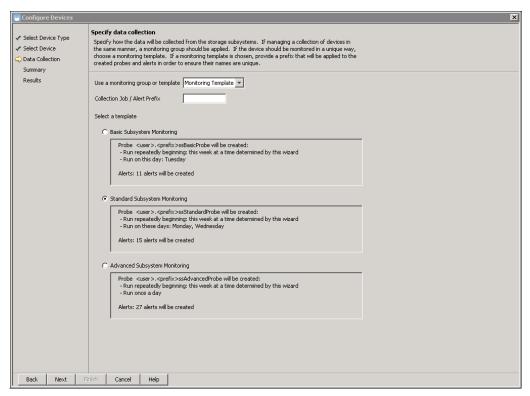


Figure 7-17 Data collection with Monitoring Templates

Tip: When you enter the string for the Collect Job / Alert Prefix field, end the string with some kind of separation character for example the underscore (_) character.

5. Summary:

The summary panel (Figure 7-18) shows the details about the probe job that the device or devices will be assigned to, and also the alerts.

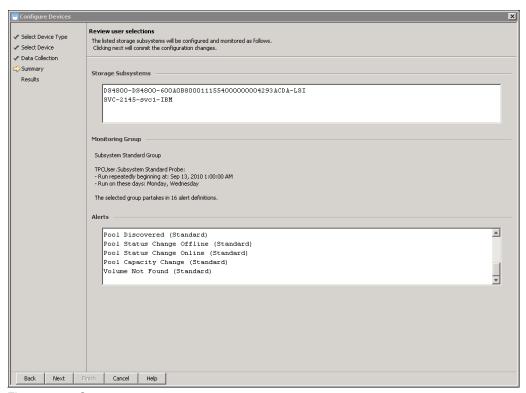


Figure 7-18 Summary

6. Results:

On the results panel (Figure 7-19), you see whether the process was successful in adding the devices to the groups and jobs, and whether a probe for the devices was started.

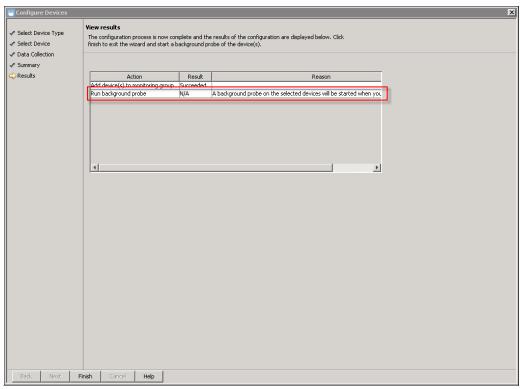


Figure 7-19 Configure Devices Wizard results page

This probe is a special job. It is not the job that is defined for your device, because that could also run the probe for other devices defined in the same job. If that probe was running, starting it again results in errors, because you cannot run two probes for a single device from a single server at the same time. This special job is the *CLI and Event Driven Job*.

First, the dialog box in Figure 7-20 opens, until you select the **Do not display this dialog again** check box.

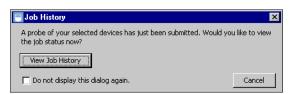


Figure 7-20 Job History

If you click **View Job History**, Tivoli Storage Productivity Center opens the Job Management panel (Figure 7-21).

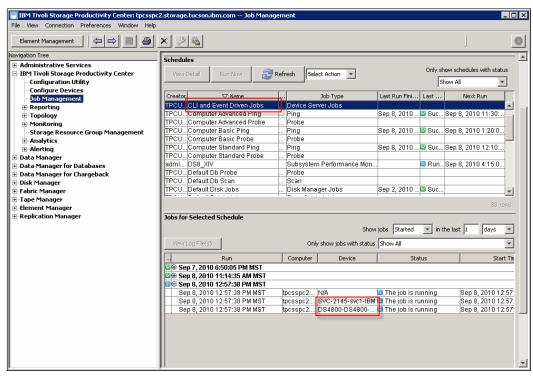


Figure 7-21 Probe started through CLI and Event Driven Jobs

7.3.2 Adding or configuring an IBM DS8000

Adding a DS8000 to Tivoli Storage Productivity Center can be easier with the use of the NAPI. Provide the following information (shown in Figure 7-22) to use the NAPI.

- IP address of the primary HMC
- Username
- Password
- Optional: The IP address of a secondary HMC

Device Type IBM DS8000	₹		
HMC Address:	9.11.99.150		
HMC2 Address (Optional) :			
Username :	tpcsspc2		
Password:	******		
Add Clear			

Figure 7-22 Adding a DS8000

7.3.3 Adding or configuring an IBM SAN Volume Controller (SVC) or Storwize V7000

After you read the planning and considerations to add an SVC or Storwize V7000 to Tivoli Storage Productivity CenterV4.2, we show you the central panel and the information you must provide.

The possible scenarios for adding or configuring an SVC or Storwize V7000 and the information that is required are outlined in Table 7-5.

Table 7-5 SVC and Storwize V7000 scenarios

Scenario	Use existing uploaded key	Upload new keys
SVC Version 4 Remember: SVC4 does not associate an SSH key with a user so you do not need to provide a user name.	Provide information: ► Select SVC version ► Enter IP address ► Private SSH key ► Enter passphrase if the key is protected	Provide information: ► Select SVC version ► IP address ► Enter admin user name Tivoli Storage Productivity Center needs the SVC admin user ID to be able to upload the key ► Enter admin password ► Enter private SSH key ► Enter passphrase if the key is protected
SVC or Storwize V7000 Version 5+	Provide information: ► Select SVC or Storwize V7000 version ► Enter IP address ► Enter private SSH key ► Enter passphrase, if the key is protected	Provide information: ► Select SVC or Storwize

Note the following additional information:

- ► Make sure that you have selected the right SVC or Storwize V7000 version, Tivoli Storage Productivity Center sometimes does not do this automatically for you.
- ► Tivoli Storage Productivity Center does not store any of the following information in its database because the information is needed only for adding the SVC or Storwize V7000 not for communication with it later:
 - Admin username
 - Admin password
 - Username

Remember, any SSH session to an SVC or Storwize V7000 is not opened using a real user name. A session is always initiated with the user name string **admin** and the correct key, so the user name is required only when you upload a new key for SVC or Storwize V7000 5+.

- You do not need to select a user by using the Select User button, you can also type the name of the user.
- Avoid using admin or superuser in the Username field, because this can overwrite and lock out others when you upload new keys. Tivoli Storage Productivity Center also does not show these user IDs in the list of users when you click **Select User**.
- ▶ If you want Tivoli Storage Productivity Center to also create a new SVC or Storwize V7000 user for you and associate the key with that user, enter a name of a user that does not yet exist.
- ▶ When Tivoli Storage Productivity Center is asking for the private key, it uses that key to generate a new public key from that private key. Therefore, you provide only that one key file name.
- ► The key file must be stored on the system that the GUI runs on, so if you are remotely connected to the Tivoli Storage Productivity Center server, the browse function cannot select a key that is already stored on the Tivoli Storage Productivity Center server.
- When you run the GUI on a separate computer than the Tivoli Storage Productivity Center server, the Tivoli Storage Productivity Center GUI uploads the key to the Tivoli Storage Productivity Center server that is stored in the key file directory:
 - ...\TPC\device\cert
- ► If a key with the same file name is already stored on the server, Tivoli Storage Productivity Center appends a number to the key file name.
- The key file tpc_svc.pem that is included with Tivoli Storage Productivity Center will be in separate locations depending on whether you are currently installing the product or you have already installed Tivoli Storage Productivity Center. If you are currently installing Tivoli Storage Productivity Center and running the migration, the key file will be stored on the CD or the directory from where you are installing. When Tivoli Storage Productivity Center is already running, it is in the Tivoli Storage Productivity Center server directory tree:
 - ...\TPC\device\conf\
- ► If you delete an SVC or Storwize V7000 from Tivoli Storage Productivity Center, the key file will not be deleted from the key file directory.

Figure 7-23 shows an example of adding an SVC Version 5 to Tivoli Storage Productivity Center.

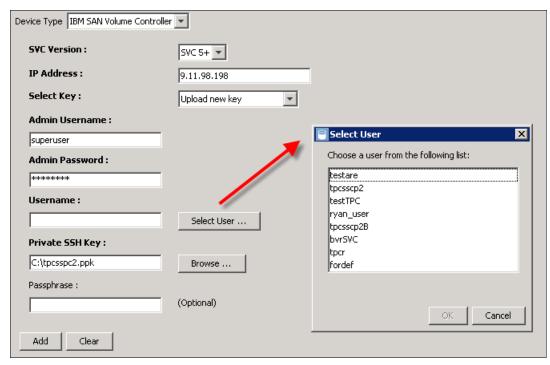


Figure 7-23 Adding an SVC to Tivoli Storage Productivity Center

7.3.4 Adding or configuring an IBM XIV

Adding an IBM XIV to Tivoli Storage Productivity Center is even easier than adding a DS8000, because it requires less information to be provided (Figure 7-24). The following information is required:

- ► IP address
- ▶ Username
- Password

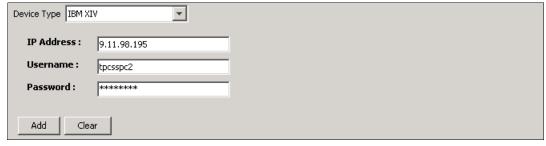


Figure 7-24 Adding a XIV

7.3.5 Adding or configuring a CIMOM

The information that is required for adding a CIMOM to Tivoli Storage Productivity Center V4.2 is the same as in the previous version, but the panel differs slightly. The most noticeable change is that you do not need to type the password a second time.

Depending on the type of the device you are adding, the panel can have different headlines, but in general it looks like the panel Figure 7-25 for storage subsystems and tape libraries.

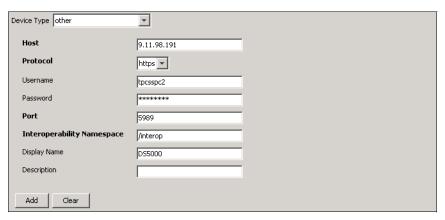


Figure 7-25 Adding a CIMOM

7.4 Adding or configuring fabrics and switches

In the same way you use the Configure Devices wizard to add storage subsystems, you can also add fabrics and switches to Tivoli Storage Productivity Center.

Figure 7-10 on page 243 shows that you can select Fabrics/Switches to configure. In 7.3, "Using the Configure Devices wizard" on page 241, we described adding and configuring storage subsystems. Now, we cover the Fabrics/Switches selection.

For each device, we provide a table containing high-level requirements. This information can help you to understand which type of data source you need to set up. More detailed information is available in Appendix A Supported Fabric Agent Types in the *Tivoli Storage Productivity Center V4.2 User's Guide*, SC27-2338 and should be used as a planning tool.

As you are adding a switch, always read the additional information at the top of the window, as shown in Figure 7-26.

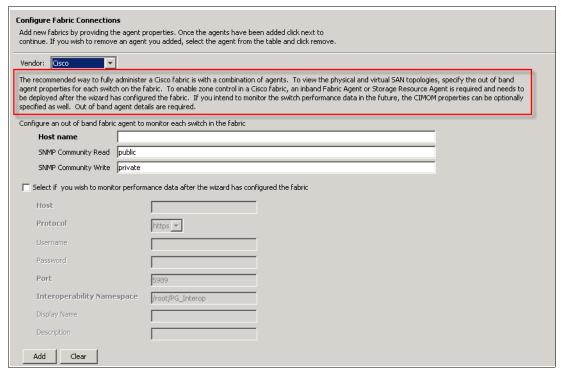


Figure 7-26 Additional details of data sources when you add a switch

We do not describe the In-band Fabric agent (available with earlier versions of Tivoli Storage Productivity Center) because it has been replaced by the Storage Resource agent. The In-band Fabric agent is still supported and you can find the details of when to use that agent in Appendix A Supported Fabric Agent Types in the *Tivoli Storage Productivity Center V4.2 Users Guide*, SC27-2338.

Besides running the Configure Devices wizard, you should not forget to set up the switches to send SNMP traps to Tivoli Storage Productivity Center.

7.4.1 Adding or configuring a Brocade Fabric or switch

When you add a Brocade Fabric/Switch, you must determine which information you want to see in Tivoli Storage Productivity Center and what functions that you want to invoke. Table 7-6 shows the high-level requirement, and which data sources can be used.

Table 7-6 High-level requirements for Brocade

Requirements	Brocade
Topology	 Recommended: CIMOM agent Also supported: out-of-band Fabric agent or Storage Resource agent
Zone control	CIMOM agent Note: Zone control with zone aliases is supported.
Switch Performance	CIMOM agent
Events	Manually add Tivoli Storage Productivity Center as an SNMP receiver in the switch configuration. This is done on the switch and not with the Configure Devices wizard.

Because the CIMOM is the preferred data source, the wizard preselects this option (Figure 7-27).

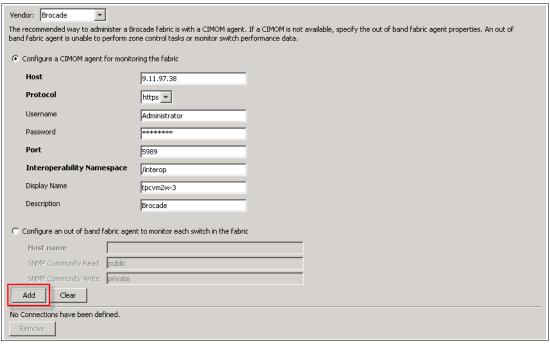


Figure 7-27 Add Brocade switches

7.4.2 Adding or configuring a McDATA Fabric or switch

When you add a McDATA Fabric or switch, you must determine which information you want to see in Tivoli Storage Productivity Center and the functions that you want to invoke. Table 7-7 shows the high level requirements.

Table 7-7 High level requirements for McDATA

Requirements	McDATA	
Topology	 Recommended: CIMOM agent Also supported: Out-of-band Fabric agent or Storage Resource agent 	
Zone control	 Recommended: CIMOM agent Also supported: Storage Resource agent Note: Zone control with zone aliases is not supported. 	
Switch Performance	CIMOM agent	
Events	Manually add Tivoli Storage Productivity Center as an SNMP receiver in the switch configuration. This is done on the switch and not with the Configure Devices wizard.	

Because the CIMOM is the preferred data source, the wizard preselects this option, as shown in Figure 7-28.

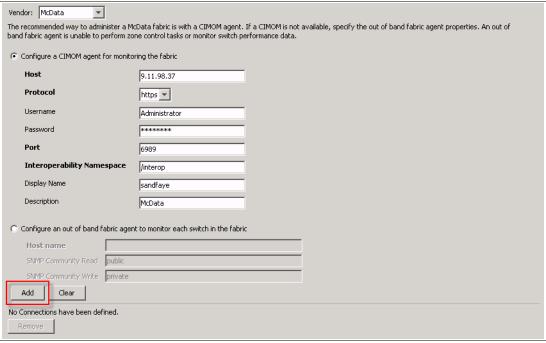


Figure 7-28 Add McDATA switch

7.4.3 Adding or configuring a Cisco Fabric or switch

When you add a Cisco Fabric or switch you need to determine which information you want to see in Tivoli Storage Productivity Center and the functions that you want to invoke. Table 7-8 shows high-level requirements.

Table 7-8 High-level requirements for Cisco

Requirements	Cisco
Topology	Out-of-band fabric agent or Storage Resource agent
Zone control	Storage Resource agent required in each VSAN Note: Zone Control with Zone Aliases is not supported
Switch Performance	CIMOM agent
Events	Manually add Tivoli Storage Productivity Center as an SNMP receiver in the switch configuration. This is done on the switch and not with the Configure Devices wizard.

Although Storage Resource agents are also supported for getting topology information, because you are in the process of adding a switch, the wizard preselects the out of band Fabric agent, as shown in Figure 7-29.

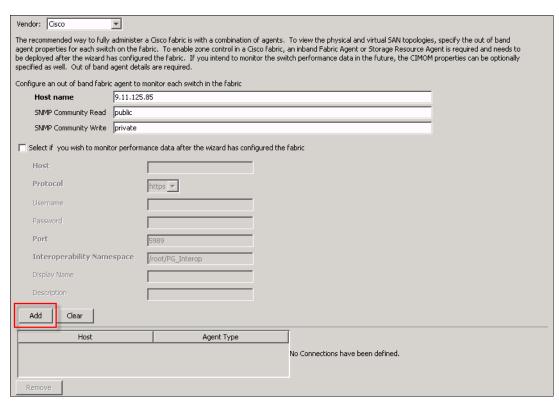


Figure 7-29 Add Cisco Switch

7.4.4 Adding or configuring a QLogic Fabric or switch

When you add a QLogic Fabric or switch you need to determine which information you want to see in Tivoli Storage Productivity Center and the functions that you want to invoke. Table 7-9 lists the high level requirements.

Table 7-9 High level requirements for QLogic

Requirements	QLogic
Topology	Out-of-band Fabric agent or Storage Resource agent
Zone control	Storage Resource agent Note: Zone control with zone aliases is not supported
Switch Performance	CIMOM agent Note: Switch performance monitoring using the QLogic CIMOM is supported for only certain models. For support matrix see "Tivoli Storage Productivity Center support matrix" on page 792.
Events	Manually add Tivoli Storage Productivity Center as an SNMP receiver in the switch configuration. This is done on the switch and not with the Configure Devices wizard.

Although Storage Resource agents are also supported for getting topology information, because you are in the process of adding a switch, the wizard preselects this option, as shown in Figure 7-30.

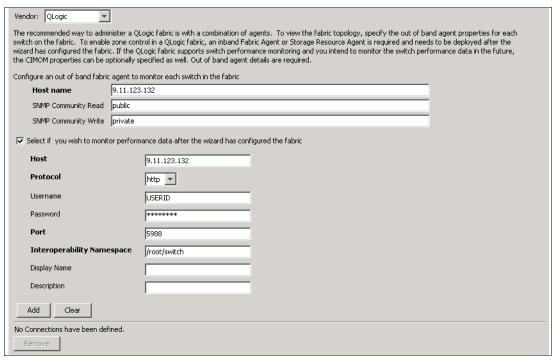


Figure 7-30 Add QLogic Switch

7.4.5 Adding or configuring a mixed Brocade/McDATA Fabric

This panel is essentially the same as the separate panels for Brocade and McDATA. This panel helps you to differentiate what kind of fabric you are adding.

If you use Brocade DCFM 10.4 to mange your environment, be sure to use the embedded CIMOM in DCFM. With that CIMOM, DCFM can manage pure Brocade fabrics (FOS), pure McDATA fabrics (EOS), or mixed fabrics (FOS and EOS). Using that CIM agent will enable you to add a single CIMOM to Tivoli Storage Productivity Center. Usually there is a fail-over concept for DCFM server (even if it is manual) so by that you might not need more than one CIM agent.

Support: Brocade Network Advisor has replaced DCFM and is supported starting with Tivoli Storage Productivity Center 4.2.1 fix pack 4.

See Chapter 14, "Fabric enhancements" on page 555 for additional information.

If you are not using Brocade DCFM 10.4 or Brocade Network Advisor, you will need to use separate proxy SMI agents for both types of switches and add both to Tivoli Storage Productivity Center to fully manage the heterogenous fabric.

7.5 Other enhancements and changes

This section provides information we found helpful related to the NAPI from our lab.

7.5.1 Help panels

In general, the enhanced help windows are useful throughout the tasks.

7.5.2 Deleting a device

When you delete a device, the performance monitoring job is deleted also, so there is no need for a clean up at this time.

7.5.3 Alert log

If you click **Delete all** within the **Alert Log** panel, Tivoli Storage Productivity Center deletes all alerts, not only the alerts shown on the currently opened panel.

7.5.4 Job history

In Tivoli Storage Productivity Center V4.2, individual job runs are no longer shown as a subtree that you can expand under the job name Navigation Tree entry. You can still right-click a job, and select **Job History**, which opens the Job Management panel, and highlight the job that you have selected.

7.5.5 Job Management

Tivoli Storage Productivity Center 4.2 includes a new Job Management panel, which we describe in Chapter 13, "Job Management panel" on page 537.

Storage Resource Agent

IBM Tivoli Storage Productivity Center now includes full host monitoring functionality with the light-weight Storage Resource Agent (SRA), including disk, file system, database, directory, file, and virtual environment information.

In this chapter, we explain the Storage Resource Agent (SRA) function, deployment, and data collection.

8.1 Overview

The Storage Resource Agent was introduced in Tivoli Storage Productivity Center V4.1 as a lightweight agent to collect host disk and file system information. In Tivoli Storage Productivity Center V4.2, this function is enhanced to include the following items:

- File system scan
- Database monitoring
- N-Series support including automatic discovery and manual entry
- ► Fabric management:
 - Collect topology information
 - Collect zone, zone set information
 - Perform zone control
 - Perform agent assignment
- ▶ Tivoli Storage Manager Integration
- Batch reporting changes
- ► Path planner support
- Data Sources Panel improvements
- ► IBM Tivoli Monitoring Integration

Storage Resource Agents either can be deployed remotely from the Tivoli Storage Productivity Center GUI or can be locally installed on the individual host computers.

8.2 Storage Resource Agent requirements

Storage Resource Agents are supported on most open systems host platforms. For a list of supported platforms, go to the Tivoli Storage Productivity Center support site:

http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition

Click the **Documentation** link, and enter Platform Support: Agents, Servers and GUI in the Search support box. Click the link to the Tivoli Storage Productivity Center V4.2 document.

Support: The Storage Resource Agent is not supported on SAN FS, NetWare or HP-UX on PA-RISC. Use the legacy Data agents to monitor these environments. Filesystem Extension is also no longer supported with Storage Resource Agents.

8.2.1 User requirements

To create a deployment job in the Tivoli Storage Productivity Center GUI, you must be logged in with a user that has the Productivity Center Administrator role (under Role-to-Group mappings).

Before you can create deployment schedules and deploy Storage Resource Agents on target computers, you must have the following permissions:

► A user ID that has administrative rights on those computers. You enter this ID when creating a deployment schedule. Tivoli Storage Productivity Center uses this ID to log in to the target computers and install and configure the necessary runtime files for the agents.

The user under which a Storage Resource Agent (daemon or non-daemon) runs must have the following authorities on the target computers:

- On Linux or UNIX, the user must have root authority. By default, an agent runs under the user 'root'.
- On Windows, the user must have Administrator authority and be a member of the Administrators group. By default, a Storage Resource Agent runs under the 'Local System' account.
- ► Storage Resource Agents do not collect information about orphan zones. An *orphan zone* is a zone that does not belong to at least one zoneset.
- ► During deployment, the server communicates with the target computer using one of the following protocols:
 - Windows server message block protocol (SMB)
 - Secure Shell (SSH) protocol
 - Remote execution (REXEC) protocol
 - Remote shell (RSH) protocol
- ▶ If RSH is configured to use a user ID and password, the connection fails. To successfully connect to a system using RSH, you must set up the .rhosts file (in the home directory of the account). RSH must be configured to accept a login from the system that is running your application.
- ► If you want to install a Storage Resource Agent or Data agent on Windows targets, the **Enable NetBIOS over TCP/IP** option must be enabled in the Control Panel settings for the computer's network connections properties.
 - To set this option on Windows 2003, click Start → Settings → Control Panel →
 Network and Dial-Up Connections → <a_connection> → Properties → Internet
 Protocol (TCP/IP) → Advanced Tivoli Storage Productivity Center WINS →
 Enable NetBIOS over TCP/IP.
 - To set this option on Windows 2008, click Start → Control Panel → Network and Sharing Center → Manage Network connections → Double Click respective Interface → Properties → Highlight the respective protocol TCP/IPv4 or TCP/IPv6 → Click Properties → Advanced → Choose WINS Tab → Select appropriate NetBIOS over TCP/IP setting.

See the documentation for your firewall to determine which ports are blocked for inbound requests.

Attention: On Windows 2008, make sure to turn off the Windows firewall before deploying the Storage Resource Agent. If you do not turn off the firewall, the deployment will fail.

8.2.2 Platform dependencies

The Storage Resource Agent is now supported on the following platforms:

- ► HPUX 11iv3 on Itanium
- ► Solaris 9, 10

The following dependency items must be installed before you install the Storage Resource Agent:

► HP-UX 11iv3 on Itanium: Requires HP libc cumulative patch PHC_34275. The patch can be downloaded from this website:

```
https://www1.itrc.hp.com/service/home/home.do
```

► Solaris 9: Requires 111711 and 111712 patches. The patches can be downloaded from the following Oracle website:

```
http://sunsolve.sun.com/patchfinder/
```

- AIX 5.3: Minimum maintenance level required is Technology Level 5300-11 Service Pack 4.
- AIX 6.1: Minimum maintenance level required is Technology Level 6100-04 Service Pack 5.
- ► Red Hat Linux 5: Requires compat-libstdc++-296-2.96-138.i386.rpm or later

Dependencies: See the detailed platform dependencies at the following address:

http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=/com.ibm.tpc V42.doc/fqz0 r sw requirements os.html

8.2.3 Communication requirements and types

The Tivoli Storage Productivity Center server uses the following protocols to communicate with the target host computer during deployment. At least one of these protocols must be enabled for the deployment to succeed:

- Windows server message block protocol (SMB) and Common Internet File System (CIFS)
- Secure Shell protocol (SSH)
- ► Remote execution protocol (REXEC)
- Remote shell protocol (RSH)

At deployment time, you have the option to use either a daemon or a non-daemon service:

- ▶ If you deploy the Storage Resource Agent as a daemon service, that service runs in the background on the monitored computer and listens for requests from the Tivoli Storage Productivity Center server. Connectivity between the server and agent is established using SSL. The server and agent have their respective certificates and no additional information is required besides those certificates and the security that is provided by the SSL protocol.
- ▶ If you deploy the Storage Resource Agent as a non-daemon service (service on-demand), it runs as a stand-alone executable file on the monitored computer. Communication from the server to the agent uses the same protocol that was used during the deployment of the agent. Communication from the agent to the server uses SSL.

Tip: Use the daemon communication mode, which allows the Tivoli Storage Productivity Center to collect data from the Storage Resource Agents as needed. Also, you do not need to update the authentication information if the server login information changes.

For non-daemon agents, the Tivoli Storage Productivity Center server requires the correct authentication information (user ID, password, port, certificate location, or passphrase) for monitored computers each time it communicates with the Storage Resource Agents running on those computers.

If the authentication information changes for a host computer on which a non-daemon Storage Resource Agent is deployed, the authentication information for that agent must be updated using the Change Authentication or Update Storage Resource Agent List functions on the Data/Storage Resource Agents panel, which is located under the Administrative Services → Data Sources node in the navigation tree.

8.3 Storage Resource Agent installation methods

In this section we describe the various ways to install the Storage Resource Agent.

8.3.1 Local graphical installer

In Tivoli Storage Productivity Center V4.2, users are able to install the Storage Resource Agent from the local graphical installer. This installer is only enabled with installing the Tivoli Storage Productivity Center server or adding a Storage Resource Agent to the Tivoli Storage Productivity Center server computer.

When a Storage Resource Agent is installed using the local graphical installer, it needs to be uninstalled with the local graphical uninstaller.

When the user attempts to uninstall a Storage Resource Agent that was installed using the graphical installer from the Tivoli Storage Productivity Center GUI, the request will be denied.

Use either the remote deployment from the Tivoli Storage Productivity Center server or the local CLI installation.

8.3.2 Storage Resource Agent deployment from GUI

Storage Resource Agents can be remotely deployed from the Tivoli Storage Productivity Center GUI.

Complete the following steps:

 To view the currently installed Storage Resource Agents, navigate to the Administrative Services → Data Sources → Data/Storage Resource Agents (Figure 8-1).

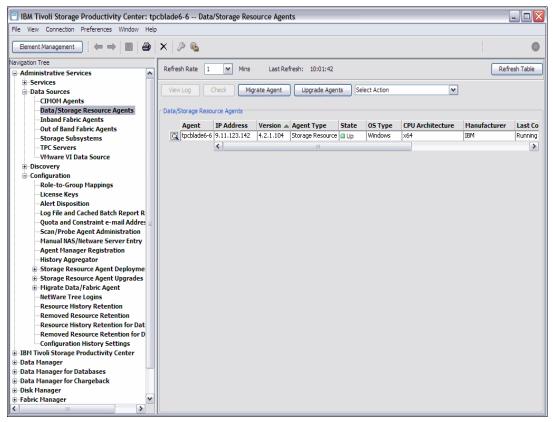


Figure 8-1 Data/Storage Resource Agents panel

2. To create a new deployment job, select **Add Storage Resource Agents** under the Select Action drop-down menu (Figure 8-2).

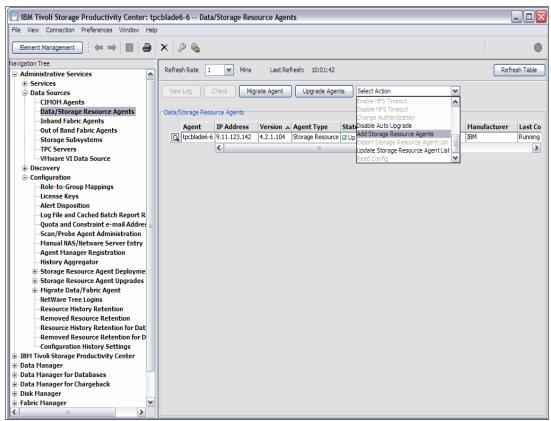


Figure 8-2 Add Storage Resource Agents selection

3. The Create Storage Resource Agent Deployments panel opens (Figure 8-3).

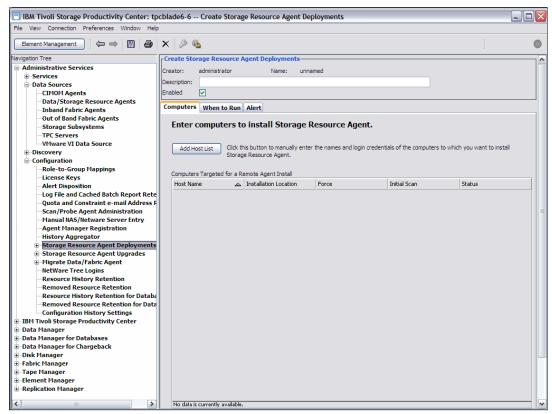


Figure 8-3 Create Storage Resource Agent Deployments

4. In this panel, click **Add Host List** to add information about the Storage Resource Agents you would like to deploy.

The Login Information panel opens (Figure 8-4).

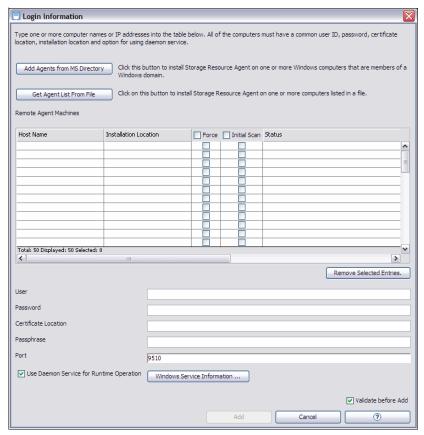


Figure 8-4 Storage Resource Agent Login Information panel

5. In this panel (Figure 8-4), enter the host name and installation location for each of the Storage Resource Agents that you want to install. Each of these systems must have a common user ID and password, which you enter on the lower half of the panel (Figure 8-5).

Tip: If you have separate user IDs and passwords, you can launch the Login Information panel once for each user ID and password combination.

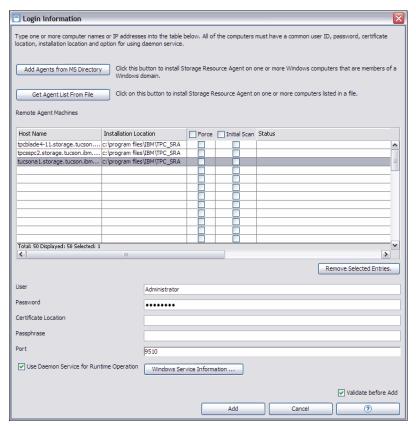


Figure 8-5 Storage Resource Agent Login Information panel

By default, Storage Resource Agents on Windows run under the Local System account. Keep this default option.

If you want to change this option to a specific user, click **Windows Service Information**. In the dialog box, specify a custom user and password (Figure 8-6). This user can be an existing local or domain user, or you can specify a new local user and that user will be created.

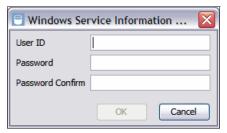


Figure 8-6 Storage Resource Agent Windows Service Information panel

You can also install Storage Resource Agents on computers in a Windows Active Directory Domain by clicking **Add Agents from MS Directory** on the Login Information panel. A panel that requests domain controller information opens (Figure 8-7).



Figure 8-7 Storage Resource Agent Domain Login Information

After you enter the domain controller information, Tivoli Storage Productivity Center authenticates to the active directory domain and displays all the systems available in that domain (Figure 8-8).

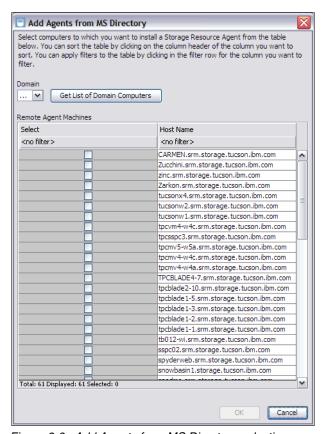


Figure 8-8 Add Agents from MS Directory selection

6. Select the agents to which you want to deploy Storage Resource Agents. After you do so, they are automatically added to the Storage Resource Agent Login Information panel and validated to ensure that the proper installation prerequisites are met.

You are returned to the Storage Resource Agent deployments panel, which lists each Storage Resource Agent to be installed (Figure 8-9).

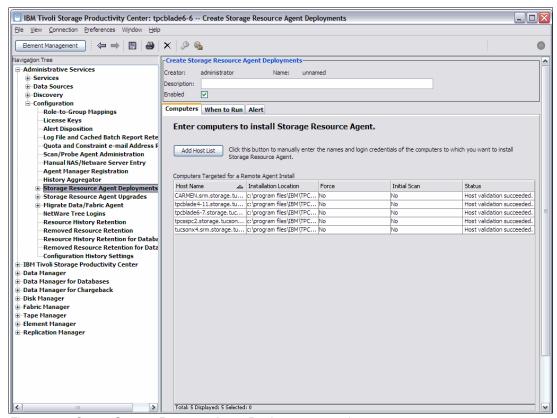


Figure 8-9 Create Storage Resource Agent Deployments panel

7. If you want to schedule the Storage Resource Agent deployment job to run at a later time, click the **When to Run** tab and choose when you would like to job to run (Figure 8-10). By default, the job will be run immediately after you save the job.

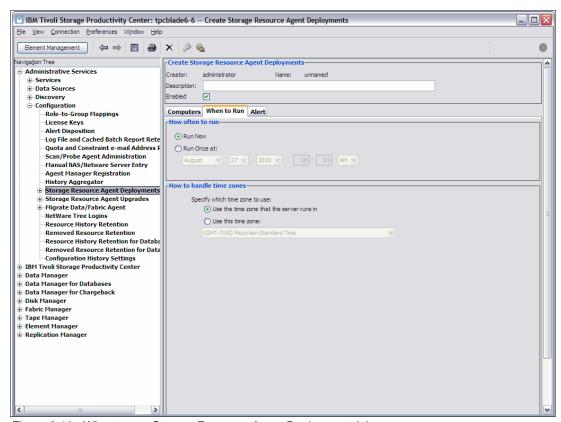


Figure 8-10 When to run Storage Resource Agent Deployment Job

If the Storage Resource Agent deployment job fails, you can set Tivoli Storage Productivity Center to send alerts under the Alert panel (Figure 8-11).

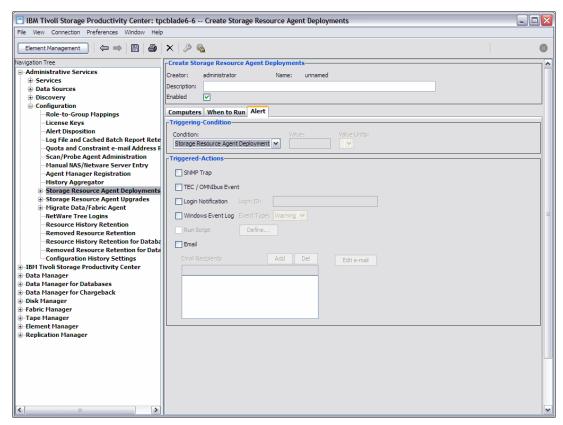


Figure 8-11 Storage Resource Agent deployment alert configuration

 After you verify the Storage Resource Agent deployment information, click File → Save to save and run the Storage Resource Agent deployment job. You are prompted for a Storage Resource Agent Deployment name (Figure 8-12). Add a descriptive name and click OK.



Figure 8-12 Specify Storage Resource Agent deployment name

9. Confirm that you want to view the job status information (Figure 8-13).

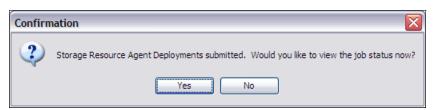


Figure 8-13 Storage Resource Agent deployment confirmation

10. Click **Yes**. The Job Management panel for the Storage Resource Agent deployment job opens (Figure 8-14).

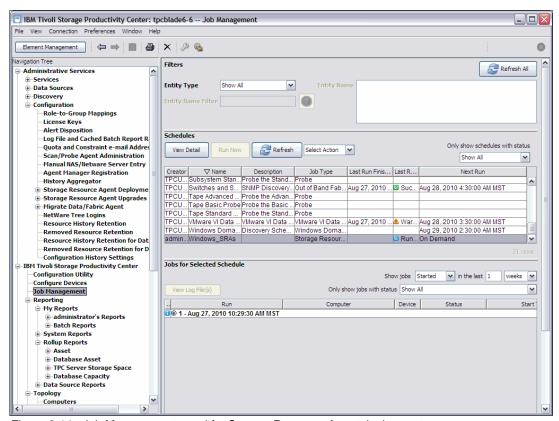


Figure 8-14 Job Management panel for Storage Resource Agent deployment

11. Within the Job Management panel, click **Refresh All** to update the job status. After a few minutes, the job completes and you can view the job logs (Figure 8-15).

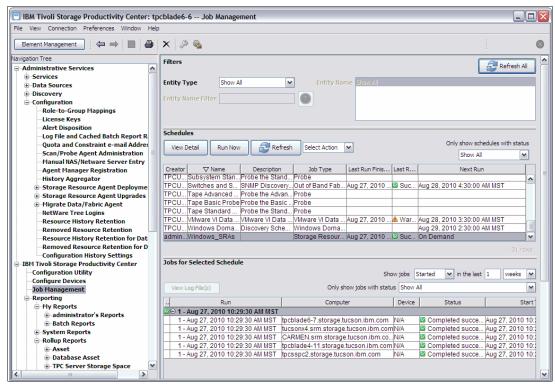


Figure 8-15 Job Management panel for Storage Resource Agent Deployment Job

After the deployment job completes, you are able to see the Storage Resource Agents under Administrative Services \rightarrow Data Sources \rightarrow Data/Storage Resource Agents (Figure 8-16).

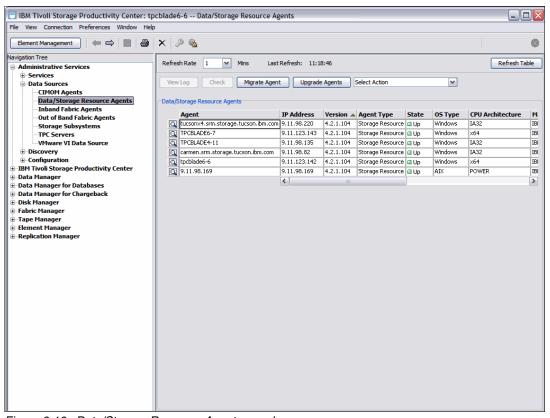


Figure 8-16 Data/Storage Resource Agents panel

8.3.3 Local or CLI installation of the Storage Resource Agent

Although we suggest to install a Storage Resource Agent by using the Tivoli Storage Productivity Center GUI, Storage Resource Agents can also be installed by using command-line interface (CLI) options.

The GUI installation method provides you with more details about the installation and any failures. For example, with the CLI installation, you will get a return code of zero for a successful installation and a nonzero return code for an unsuccessful installation. If an installation is unsuccessful, review the log files to determine the cause of the failure.

The **-force** option can be used when you have a Storage Resource Agent pointing to multiple servers. If one server has installed an agent, and another server wants to install an agent in the same location, the second server can use the **-force** option to install the agent.

Tip: When you use the CLI installation, specify an installation directory that does not end in a backslash mark (\). For example, the following line can cause the installation to fail: C:\agent\

If you run the agent as a non-daemon service (on-demand service), you must make sure that at least one protocol is valid for a successful connection from the server to the agent. Table 8-1 lists the required parameters for each protocol.

Table 8-1 Storage Resource Agent non-daemon protocols

Protocol	Description
SSH	Requires the user ID and password or user ID, certificate, and passphrase.
Windows (SMB, CIFS)	Requires the user ID and password.
REXEC	Requires the user ID and password.
RSH	Requires the user ID.

Image location

The installation images are contained in either the disk1 or the Storage Resource Agent images. The images for each platform are located on the paths listed for operating system (from Table 8-2). The images are located in the following directory:

TPC installation image location/data/sra/operating system

Table 8-2 shows the Storage Resource Agent installation images.

Table 8-2 Storage Resource Agent image path names

Installation image	Operating system
AIX	aix_power
HP-UX	hp-ux_itanium
Linux x86	linux_ix86
Linux for IBM Power Systems™	linux_power
Linux s390	linux_s390
Solaris	solaris_sparc
Windows	windows

Installing the Storage Resource Agents through the CLI

We show an example of how to install a Storage Resource Agent with daemon communication. For non-daemon installations and additional installation parameters, see the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide*, SC27-2337.

To install the Storage Resource Agents locally, complete the following steps:

- 1. Copy the agent installation image to a local file system (or insert a DVD with the installation images).
- 2. Open a Command Prompt window.

3. Navigate to the installation image path for the appropriate operating system and run the Agent.exe installation executable file. See Figure 8-17 and Table 8-3.

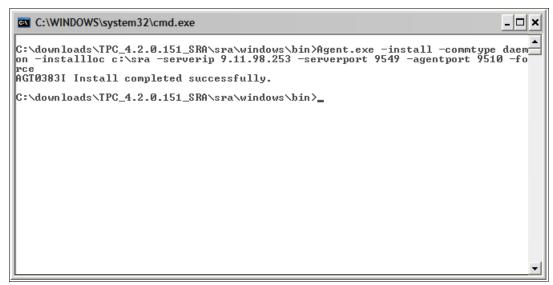


Figure 8-17 Storage Resource Agent command-line installation

Table 8-3 Storage Resource Agent CLI installation parameters

Option	Description
commtype	Use only when installing a daemon agent
installLoc	Location where the agent is installed. Enclose the directory name in quotation marks, for example: "C:\Program Files\IBM\TPC_SRA\"
serverip	IP address of the Tivoli Storage Productivity Center server
serverport	Port for the Data server. The default is 9549
agentport	If the agent is run as a daemon service, the agent port must be specified.
debug	Optional parameter for debugging purposes
duser	For daemon service; the login account that is used to start the Storage Resource Agent service on Windows
dpassword	For daemon service; password for the Windows login account that is specified by the -duser parameter

If the installation fails, see the return codes in the Tivoli Storage Productivity Center information center. Search for Return codes used by Storage Resource Agent.

8.4 Database monitoring with Storage Resource Agent

This section describes Tivoli Storage Productivity Center Data Manager for Databases, which is similar to the Tivoli Storage Productivity Center Data Manager. The similarity is in that reporting uses the data in the repository and the reports can be displayed as either tables or as graphical charts.

Data Manager for Databases can be integrated into Oracle, Microsoft SQL/Server, Sybase, and IBM DB2 Universal Database, and can extract details about database and table space usage. You can also enable monitoring for predefined events to occur and alert you to potential situations before they occur.

For example, it can notify you when an Oracle table space is reaching a critical shortage of free space or when a Sybase table is dropped. By alerting you to these and other issues related to your stored data for the databases within your environment, it enables you to prevent unnecessary system and application downtime.

In this section, we show you how to register a database and extract capacity and usage reports.

To better demonstrate all of the reporting types, see Figure 8-18. Note that not all entries are fully expanded.

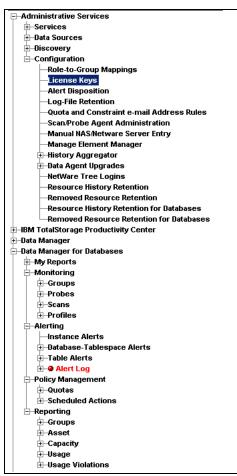


Figure 8-18 Data Manager for Databases reporting areas

8.4.1 Registering the database

Before you can create probes and scans and then extract data from our databases, first register the database to Tivoli Storage Productivity Center for Data. In these steps, we register our database instance, which is installed on one of our hosts, to Tivoli Storage Productivity Center for Data:

- 1. Select Administrative Services → Configurations → License Keys (Figure 8-19).
- 2. Click the magnifying glass icon that is located to the left of the line Tivoli Storage Productivity Center for Data Databases.

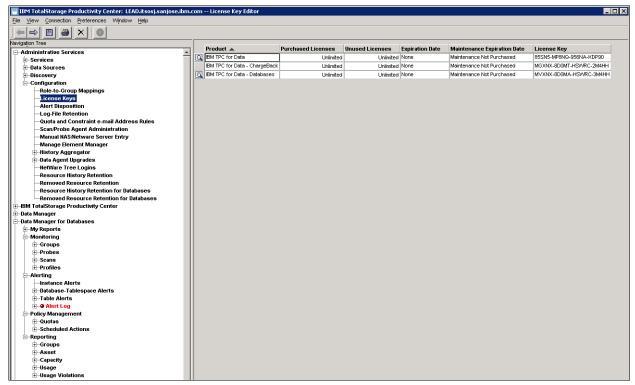


Figure 8-19 Register the database

3. In the next window (Figure 8-20), select the RDBMS Logins tab.

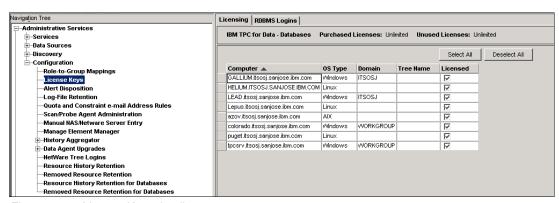


Figure 8-20 License Keys detail

4. Click Add New.

5. In the window that opens (Figure 8-21), configure the login properties for your database instance on the target server, and then click **Save**.

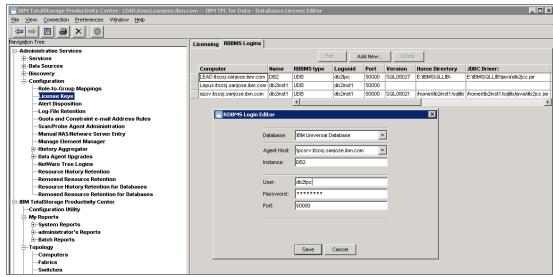


Figure 8-21 RDBMS Login Editor window

6. The success message is displayed (Figure 8-22). Click **OK**. You now see that the new database host and information are listed after the configuration.

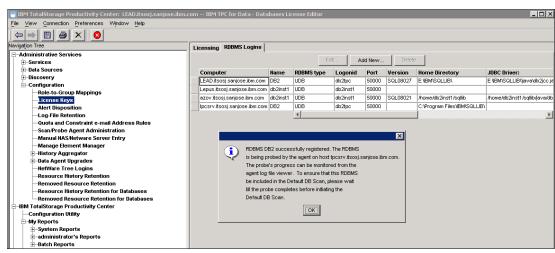


Figure 8-22 RDBMS success message

8.4.2 Setting up probes and scans

After registering the database, create a probe to extract information from the database:

 Select Data Manager for Databases → Monitoring → Probes. Right-click and select Create Probe (Figure 8-23).

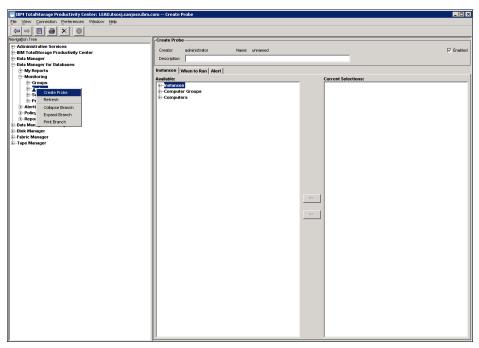


Figure 8-23 Creating a probe

 You can add Instances and use the arrows to move the instances to the Current Selections panel. Select File → Save. In the Save As pop-up window (In Figure 8-24), give the probe a name and click OK. The probe is submitted.

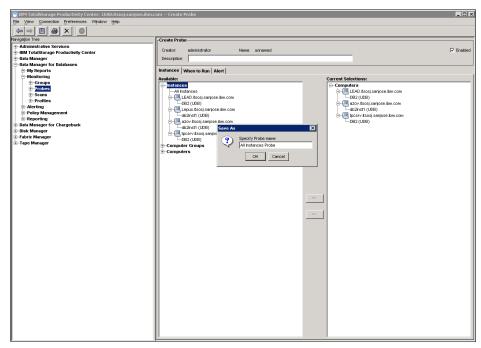


Figure 8-24 Configure and save the probe

3. Monitor the job until it completes (Figure 8-25).

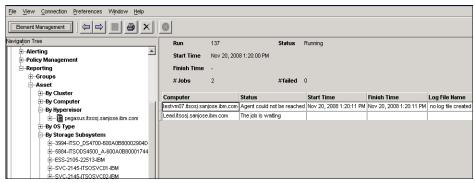


Figure 8-25 View the probe job's success

 Now create a scan to gather more detailed information about the database. Select Data Manager for Databases → Monitoring → Scans. Right-click the default scan and select Run Now (Figure 8-26).

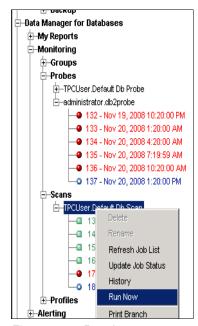


Figure 8-26 Running a scan

5. Monitor the scan job until it completes (Figure 8-27).

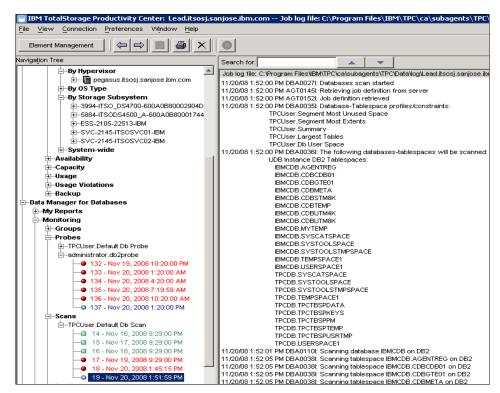


Figure 8-27 Monitor the scan

8.4.3 Database capacity reports

Capacity information for databases is available in the repository. Tivoli Storage Productivity Center for Data can report several types of detailed information that is categorized by database, instance, computer, computer group, and network-wide.

For example, to get the Total Instance Storage Capacity for a view by computer, use the following steps:

- 1. Select Data Manager for Databases \rightarrow Reporting \rightarrow Capacity \rightarrow All DBMSs \rightarrow Total Instance Storage.
- 2. Click By Computer. The Report generation window opens.
- 3. Click **Generate Report** to create the report By Computer. The report displays the capacity for each computer (Figure 8-28).

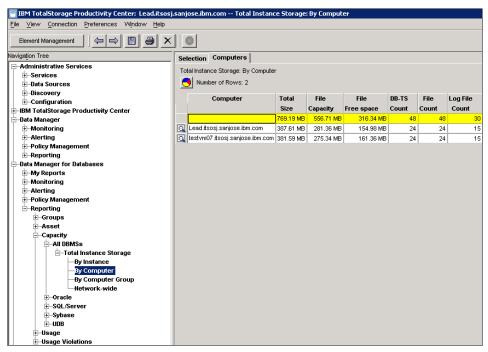


Figure 8-28 Total Instance Storage by Computer

4. To view a selected server, click the magnifying glass icon (Figure 8-29).

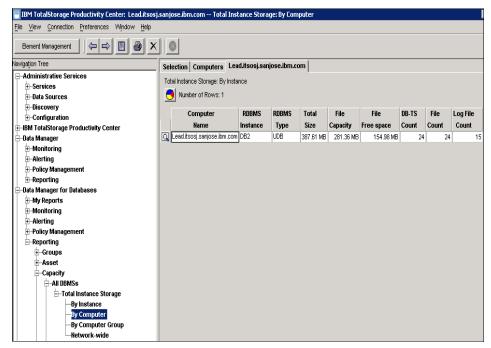


Figure 8-29 Total Instance Storage by Computer

5. Drill down on the selected computer for instance information (Figure 8-30).

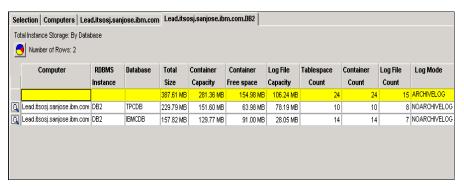


Figure 8-30 Instance Information about selected computer

6. If you click the magnifying glass icon, a list of database files is shown (Figure 8-31).

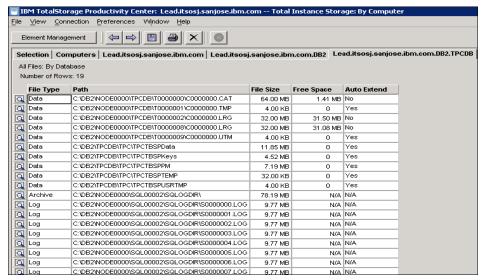


Figure 8-31 Listing of all database files

8.4.4 Database usage reports

Database usage reporting information is available in the repository. Tivoli Storage Productivity Center for Data can report on this data in several ways, such as by database, by instance, by computer, by computer group, and network-wide.

For example, to get the Total Instance Storage Capacity view *By Computer*, use the following steps:

- Select Data Manager for Databases → Reporting → Usage → All DBMSs → Databases-Tablespaces.
- 2. Click **Network-wide**. The Report generation window opens.

3. To create the report, click **Generate Report** (Figure 8-32).

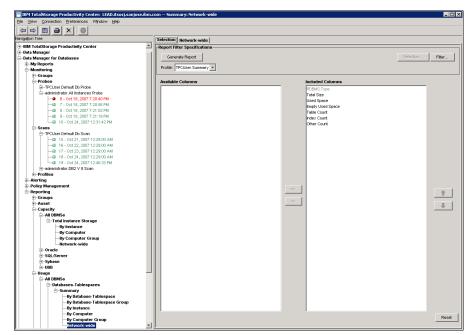


Figure 8-32 Creating a Database Usage report

The report displays a Network-wide summary report (Figure 8-33).

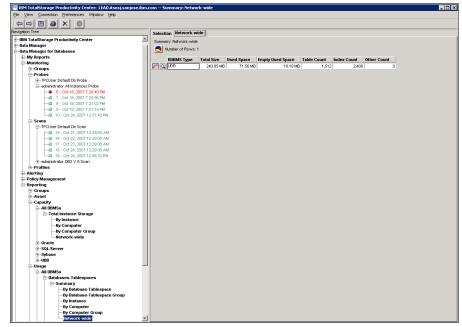


Figure 8-33 Network-wide Summary report

4. Click the line graph icon to see the Usage graph report (Figure 8-34).

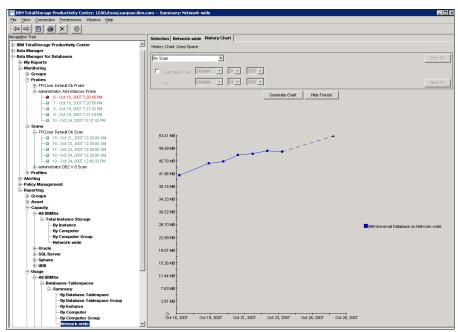


Figure 8-34 Database usage graph

8.5 IBM Power VM Virtual I/O Server (VIOS) support

In this section we show the features and installation of the Storage Resource Agent for the IBM Power VM Virtual I/O Server. Deploying an SRA on the VIOS enables you to get LUN correlation support. You can now see the back-end LUNs associated to your Power VM logical partition (LPAR) through the VIOS.

8.5.1 VIOS LUN correlation supported environments

For the LUN correlation support, you need a one to one relationship between the Virtual SCSI disk device in your Power VM logical partition (LPAR) and your storage volume attached to your Virtual I/O Server. In this setup Tivoli Storage Productivity Center can extract the storage system serial number and correlate the Virtual SCSI disk attached to the LPAR with the corresponding storage volume.

Tip: The Data Path Explorer does not work for virtual devices.

The LUN correlation support exists for the following IBM hardware:

- ► IBM System Storage DS8000 series (DS8K)
- ► IBM TotalStorage Enterprise Storage Server (ESS)
- ► IBM System Storage DS6000 series (DS6K)
- ► IBM System Storage SAN Volume Controller (SVC)
- ► IBM Storwize V7000 Midrange Disk System
- ► IBM DS3000, DS4000, DS5000 series

The LUN correlation support is currently not available for the following environments:

- IBM XIV Storage System
- ▶ non-IBM storage

Although the legacy agent is supported we suggest that you use the SRA. The SRA lets you utilize the new LUN correlation support and you can recognize virtual FC adapters. The legacy agent does not provide this support.

8.5.2 Installing the Storage Resource Agent on VIOS

Follow the remote deployment instructions in 8.3.2, "Storage Resource Agent deployment from GUI" on page 268 to successfully deploy the SRA. You need to specify an administrative ID to deploy the SRA to the VIOS. Providing the **padmin** logon is not sufficient.

8.5.3 Uninstalling the Storage Resource Agent on VIOS

Use the following command to uninstall the SRA from the VIOS:

./Agent -UNINSTALL -SERVERNAME tpcblade3-13.storage.tucson.ibm.com

Deployment using the default installation path puts the software to the /opt/IBM/TPC/agent/bin/ directory.

8.5.4 Overview of new functionality available with SRA on VIOS

If you open the Topology Viewer, you can now see the Topology of your LPAR in conjunction with the attached non-virtual physical back-end disk. As seen in Figure 8-35, the LPAR has the device /dev/hdisk2 and /dev/hdisk3 being detected as Vendor IBM and Model 2145 devices. This is the type number for SVC.

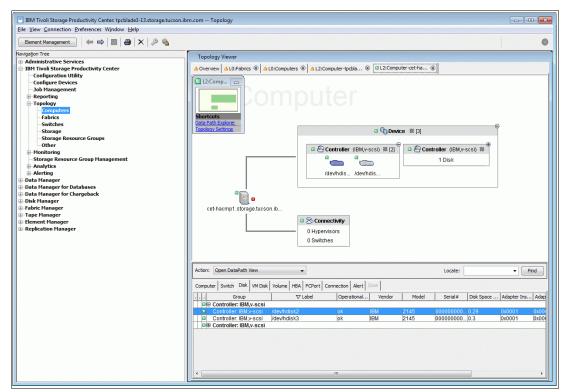


Figure 8-35 Topology viewer

Open the Datapath View by right-clicking the selected disk and you can explore the datapath from end to end as seen in Figure 8-36.

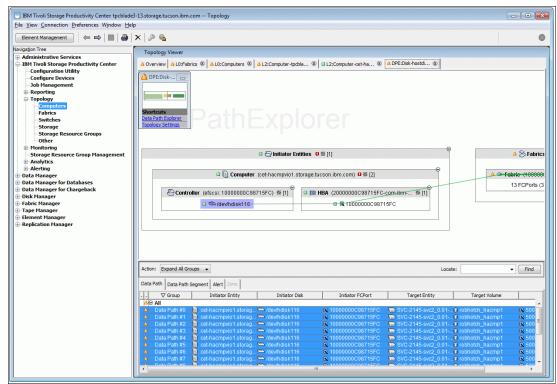


Figure 8-36 Data Path View

Also you can see detailed information about the correlated subsystem storage volume if you click the selected disk in the Data Manager as seen on Figure 8-37.

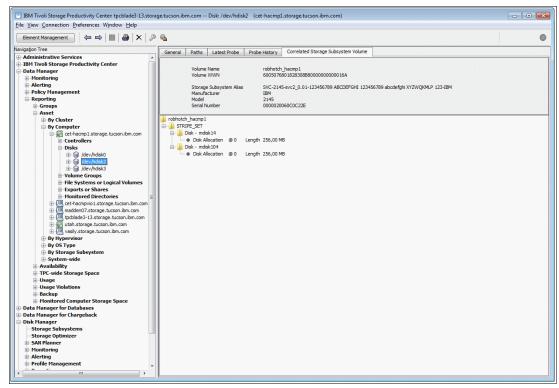


Figure 8-37 Correlated Storage Subsystem Volume

8.6 NetApp and N series monitoring

In this section we discuss Storage Resource Agents and NAS support.

8.6.1 Overview of NAS support

Tivoli Storage Productivity Center for Data supports IBM N Series and Network Appliance filers for filer discovery, probe/scan agent configuration, asset probe, file system scan, quota import and monitoring, alerting and capacity reporting.

Unlike other Storage Resource Agents, there is no agent code to install on NAS devices. Tivoli Storage Productivity Center issues SNMP queries to the NAS device to collect the summary information for the aspect of the NAS device. Tivoli Storage Productivity Center also uses a proxy agent implementation to collect more detailed information. A Storage Resource Agent is designated as the *proxy agent* responsible for collecting asset and quota information from assigned NAS devices through SNMP. Tivoli Storage Productivity Center collects the mounted file system information or shares through the Network File System (NFS) or Common Internet File System (CIFS) viewed by the system hosting the Storage Resource Agent.

After the NAS devices information is collected, the NAS devices are displayed in the Topology Viewer as computers. You can check the NAS devices information as you would computers.

The collected information is also used for the following reports (select **Reporting** \rightarrow **Asset** \rightarrow **By OS Type** \rightarrow **Network Appliance**):

- Controllers
- ▶ Disks
- ► File System or Logical Volumes
- Exports or Shares
- Monitored Directories

NAS monitoring options

The NAS device can be monitored by either a proxy Windows Storage Resource Agent or a UNIX Storage Resource Agent. Windows Storage Resource Agents are used to monitor CIFS shares on the NAS filer; UNIX Storage Resource Agents are used to monitor NFS exports. Each option has its own requirements and setup steps. In this chapter, we describe both implementation methods. See 8.6.2, "Configuring Windows Storage Resource Agent" on page 300, and 8.6.3, "Configuring the UNIX Storage Resource Agent" on page 313.

General NAS system requirements

Requirements are as follows:

- ► A NAS device must support SNMP and be enabled for queries. Check the SNMP configuration on NAS (from the FilerView by clicking **SNMP** → **Configure**) to make sure the SNMP is enabled on the NAS. By default it is enabled (Figure 8-38).
- ► Also, note the SNMP community name.

For more information about implementing and configuring Network Attached Storage (NAS) in your Tivoli Storage Productivity Center environment, see the IBM Redpaper™ publication, *IBM Tivoli Storage Productivity Center with Network Attached Storage (NAS)*, REDP-4793, at this website:

http://www.ibm.com/redbooks

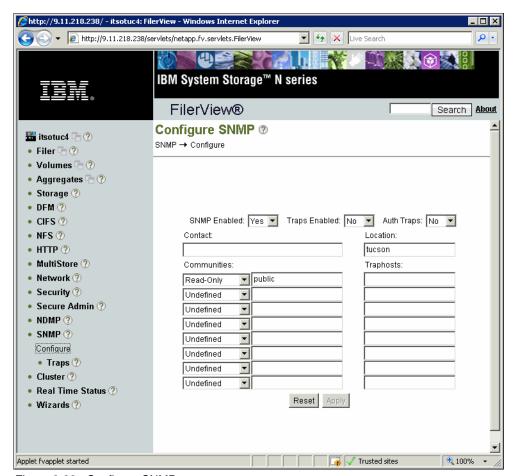


Figure 8-38 Configure SNMP

 A NAS device must supply a unique sysName, which maps to the network host name of the NAS device.

At the time of writing this book, Tivoli Storage Productivity Center V4.2 supports the IBM N series models and NetApp filers listed at the following website:

http://www-1.ibm.com/support/docview.wss?rs=597&uid=ssg1S1003019

UNIX proxy agent requirements

To monitor the NAS device through a UNIX Proxy Storage Resource Agent, the IBM N Series or NetApp filers must meet the certain criteria. NAS device must support NFS queries:

- ► The NAS must have NFS licensed and proper NFS share configured.
- ► The root file-system from the NetApp filer must be mounted on the agent where the Storage Resource Agent will be deployed. This requirement ensures that during post installation discovery, the NetApp filer will be discovered automatically.

Windows proxy agent requirements

NAS device must support Common Internet File System (CIFS) queries.

- NAS must have a licensed CIFS, and the proper CIFS share must be configured.
- ► The NAS filers within your environment must be visible to the systems where you install the agent or agents. If NAS filers are to be monitored by Windows computers, those NAS filers must be configured to be members of the same Windows domain. NAS in a Windows workgroup environment is not supported.
- ► The root file-system from the NetApp filer does not need to be mounted (on the target Windows computer), but it has to be exported. The Storage Resource Agent gets a list of computers from the browsing service for the domain that the agent machine is in.

The account that is used for scanning NAS for Windows must be a domain account that can log into both the Windows agent machine and the NAS device.

8.6.2 Configuring Windows Storage Resource Agent

This section describes the procedure to configure Tivoli Storage Productivity Center to monitor the IBM N series or NetApp filer though a Windows server.

Tip: The Windows server used as a proxy Storage Resource Agent must be a member of the Windows domain. The NAS filer also has to be added to the Windows domain.

Configuring the NAS filer to be a member of a Windows domain

You must be sure that the NAS filer is in your Windows domain. You can verify this by logging in to the Windows domain controller using a user ID with administrator privilege and clicking $\mathbf{Start} \to \mathbf{Settings} \to \mathbf{Control\ Panel} \to \mathbf{Administrative\ Tools} \to \mathbf{Active\ Directory\ Users}$ and $\mathbf{Computers}$. The panel shown in Figure 8-39 on page 300 opens. Verify that the NAS device is listed under the Computers tree.

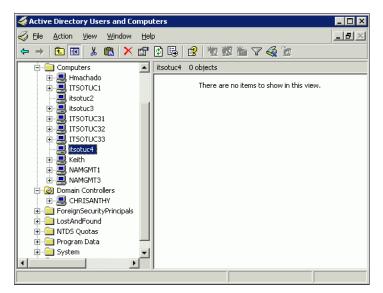


Figure 8-39 NAS is a member of a Windows domain

Configuring a CIFS share on the NAS filer

To add a new CIFS share, go to the NAS FilerView, and click CIFS \rightarrow Share \rightarrow Add, as shown in Figure 8-40.



Figure 8-40 Add a CIFS share on NAS

Mapping NAS CIFS share to Windows server, running read and write I/O

The account used for scanning NAS for Windows must be a domain account that can log into both the Windows agent machine and the NAS device. In our lab, we log in to the Windows server using such a domain account, and map the NAS CIFS share that we defined previously to the Windows server. We then do some read and write I/O on the mapped network drive (for example, copy a file to the mapped drive) to be sure that the NAS share is working correctly (Figure 8-41).

Tip: You must use a *domain user* that has *domain administrator* privileges.

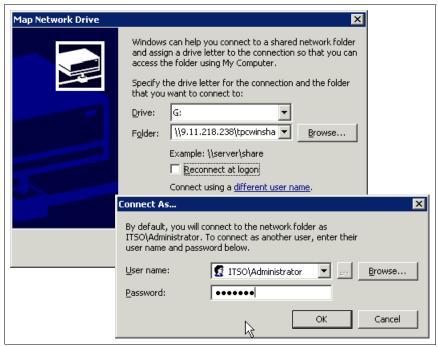


Figure 8-41 Map the NAS CIFS share

Installing Storage Resource Agent on Windows server

After installing a Storage Resource Agent on the Windows server, we use it as a proxy agent to collect information from NAS. This Windows server must be a member of the same Windows domain as the NAS or a trusted domain, it can be a domain controller or a member server.

Managing NAS devices in Tivoli Storage Productivity Center

The two ways to configure which NAS devices are managed by Tivoli Storage Productivity Center are as follows:

- ► You can run a discovery to discover the NAS devices in your environment.
- You can manually add the devices to Tivoli Storage Productivity Center.

We describe both methods.

Managing NAS devices through discovery

Important: Successful Discovery depends on the configuration being done correctly.

Complete the following steps:

- 1. Launch the IBM Tivoli Storage Productivity Center GUI.
- 2. Set a default domain login and password in Tivoli Storage Productivity Center for it to discover the NetApp devices, as follows:
 - a. Select Administrative Services → Configuration → License Keys, and double-click the TPC for Data entry (Figure 8-42).

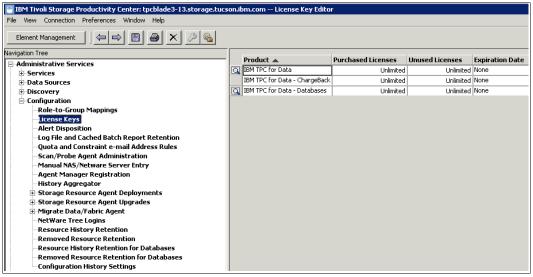


Figure 8-42 Tivoli Storage Productivity Center Licensing panel

b. Click the **Filer Logins** tab, and click **Update default login and password**. Enter the appropriate domain User ID and Password (Figure 8-43).

Tip: This user must have Domain Administrator privileges.

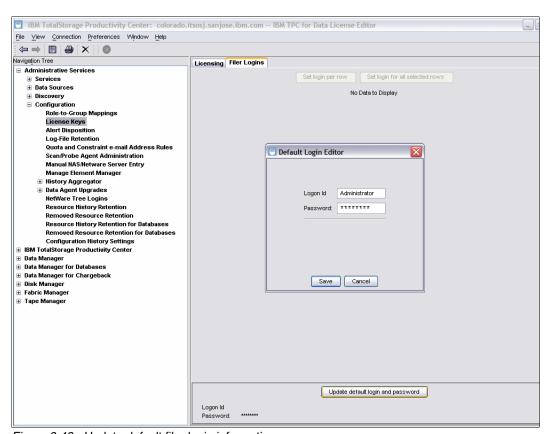


Figure 8-43 Update default filer login information

 Expand the Administrative Services in the Navigation Tree and select Administrative Services → Discovery.

SNMP: You must verify that the correct SNMP community name is defined in the Windows Domain, NAS, and SAN FS job. To verify, click **Administrative Services** → **Discovery** → **Windows Domain, NAS and SAN FS** and select the **Options** panel. Add the correct SNMP community name for the filer.

To run an NAS/NetApp Discovery job, right-click **Windows Domain**, **NAS and SAN FS** and select **Run Now** (Figure 8-44).

Tip: The discovery will discover all entities in the Windows Domain. To shorten the time taken for discovery, you can check the Skip Workstations option in the Discovery properties pane (see Figure 8-44).

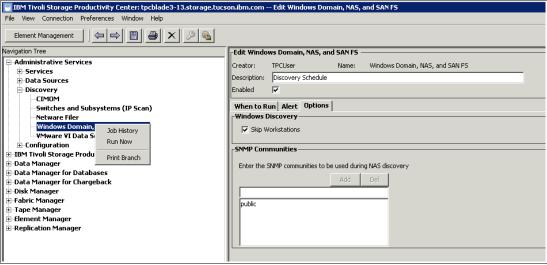


Figure 8-44 NAS Discovery job

To check the running job status, right-click **Windows Domain**, **NAS and SAN FS** and select **Update Job Status**. Wait until the discovery job finish.

4. After the discovery job completes, but before IBM Tivoli Storage Productivity Center can perform operations (Probe/Scan) against NetWare, NAS, and NetApp filers, they need to be licensed first. The reason is because, in a very large environment, the customer might not want to automatically license all the discovered NAS devices, and therefore, you have a choice of which servers to license.

Tip: For manually entered NAS, you do not need to do this step because the filer will be licensed automatically.

Select the **Administrative Services** \rightarrow **Configuration** \rightarrow **License Keys**, and double-click the **TPC for Data** entry (Figure 8-45).

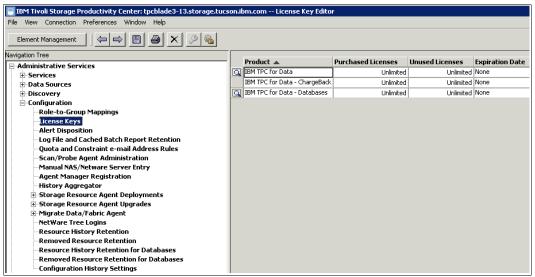


Figure 8-45 License Key panel

The Licensing tab opens. Locate the NetWare, NAS, or NetApp filer. In the Licensed column, select the associated check box for the NAS filer, and click the Disk icon on the toolbar to save changes (see Figure 8-46).

Saving: You must save the changes after you license the NAS filer before you leave this panel, otherwise Tivoli Storage Productivity Center does not save the change to its repository.

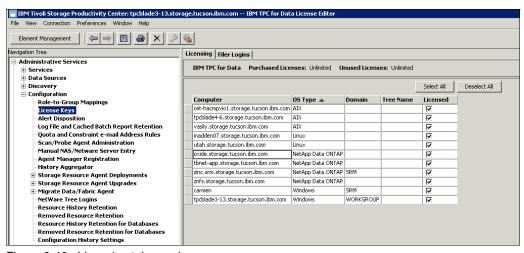


Figure 8-46 Licensing tab panel

5. Set filer login ID and password.

For Windows-attached NAS, you have to specify the login ID and password to the NAS. Click the **Filer Logins** tab, and select the NAS filer, then click the **Set login per row** button, enter the Logon Id and Password in the popped up **Filer Login Editor** panel, click **Save** to save the changes (see Figure 8-47 on page 306). This ID and password must be a Domain account that can log into both the Windows agent machine and the NAS device.

Tip: Setting filer login ID and password is not required if the filer is UNIX-attached NAS.

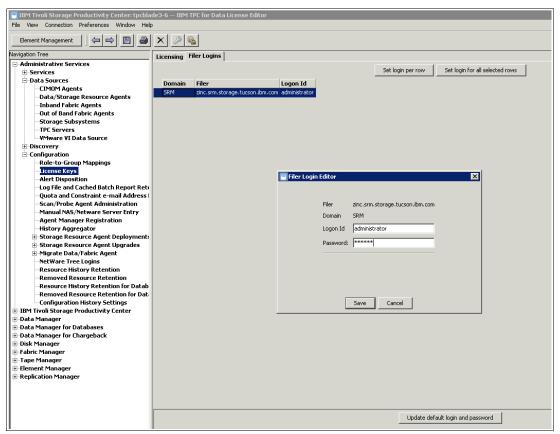


Figure 8-47 Logon ID and password for NAS filer for Windows

6. Run a discovery again:

After licensing the NAS filer and setting the login ID and password, run a discovery job again to get further information about the NAS filer. See Figure 8-44 on page 304.

Manually adding NAS device to Tivoli Storage Productivity Center

If the filer you want to monitor was not discovered, you can manually add it to Tivoli Storage Productivity Center as follows:

- 1. Launch the IBM Tivoli Storage Productivity Center GUI.
- Expand the Administrative Services in the navigation tree and select Administrative Services → Configuration → Manual NAS/Netware Server. The panel, shown in Figure 8-48, opens.

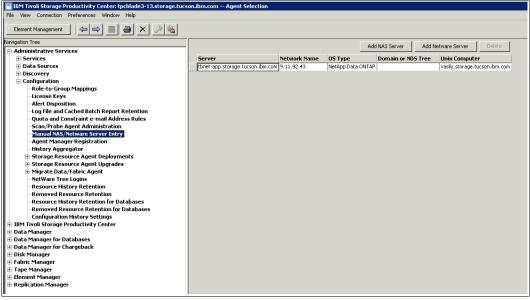


Figure 8-48 Manual NAS Server Entry Panel

3. Add NAS Server:

Click **Add NAS Server**. Enter the necessary information in the next panel (Figure 8-49) and click **OK** to continue:

Network name:

Enter the fully qualified host name or IP address of the NAS filer.

Data Manager Agent OS Type:

Select the operating system of the computer that contains the agent that will gather information about the NAS filer. In our case, we select Windows.

Accessible from:

From the drop-down list, select the agent host that you want to *discover* the NAS filer. The drop-down list will only display the following agents:

- Agents that are running under the operating system selected in the Data Manager Agent OS Type field.
- Agents that are located on Windows or UNIX computers that are accessible to the NAS filers (Storage Resource Agents are not located on the NAS filers themselves):

Windows agents are located on Windows computers within the same domain as the NAS filers.

- SNMP Community:

Enter the SNMP community name, the default is PUBLIC, which is used to get information from the NAS filer.

- Login ID and Password:

These fields are Windows only. Enter the login ID and password, which must be a Domain account that can log in to both the Windows agent machine and the NAS filer.



Figure 8-49 Add NAS Server

If all the information that you provide is correct, you see the NAS filer added to the panel (Figure 8-50).

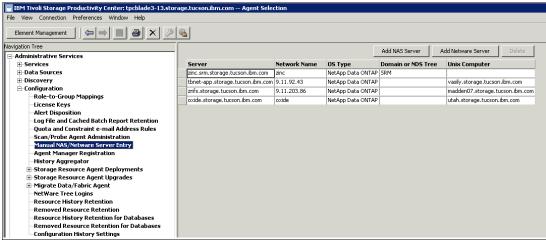


Figure 8-50 Manual Added NAS Entry

Setting the Scan/Probe Agent

After you manually added or discovered the NAS device, you need to set the Scan/Probe Agent before scanning or probing the filer. Because there is no Storage Resource Agent installed on the NAS filer itself, we need another agent running the Scan/Probe job as a proxy agent instead. After this task is done, Tivoli Storage Productivity Center will treat these devices as normal servers with attached storage. An agent can scan multiple NAS Servers, and NAS Server can be scanned by multiple agents, so we can set up parallel scans.

Remember: With this step you are assigning workload to computers. These scans create traffic on the IP network. We suggest that you select proxy agents on the same network as the NAS filer. This will maximize performance and minimize network load on other parts of your infrastructure. Scans over the network are not as performant as local scans. The scan might take extra time to run if it is doing a proxy scan on a large NAS device. For large NAS devices, use multiple proxy agents. Normally the default scan of once per day is not required, therefore, consider weekly scans.

Select Administrative Services \rightarrow Configuration \rightarrow Scan/Probe Agent Administration. Click the NAS filer that you want to define Scan/Probe agent. In our lab, we did multiple selections by pressing the Ctrl key while clicking the NAS filer entries. Click **Set agent for all selected rows**, and in the pop-up window, choose the Windows Storage Resource Agent that has the NAS filer attached (Figure 8-51).

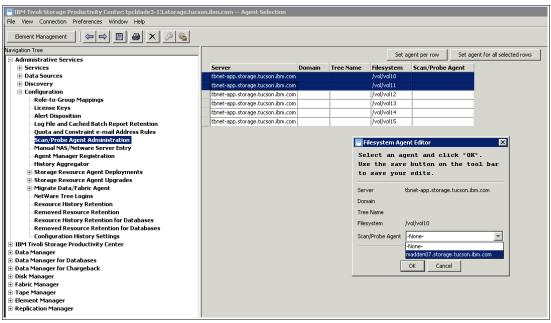


Figure 8-51 Scan/Probe agent administration

Now you can see the NAS filer file systems that have been assigned Scan/Probe Agent, make sure you click **Save** in the toolbar to save the changes (Figure 8-52).

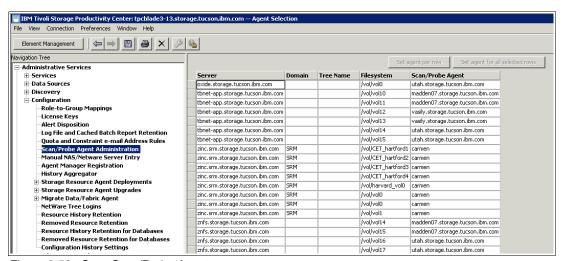


Figure 8-52 Save Scan/Probe Agent

Running a probe job

After the Scan/Probe agent has been set, a probe to collect the device hardware configuration and file system information will be run automatically. If you want to create and run an additional probe, you can do so.

In the Tivoli Storage Productivity Center GUI Navigation Tree, select **IBM Tivoli Storage Productivity Center** \rightarrow **Monitoring** \rightarrow **Probes.** Right-click and select **Create Probe**, click **What to PROBE** tab, and add the only NAS filer only, as shown in Figure 8-53.

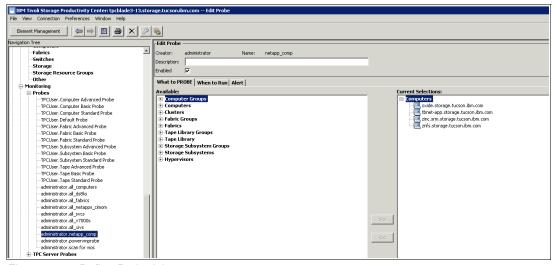


Figure 8-53 Define Probe job

Click the **When to Run** tab, select the **Run Now** button, and click **Save** to save the probe job (Figure 8-54). The probe jobs starts. You can right-click the Probe Job and select **Update Job Status** to check the running job status. Wait until the probe job finishes.

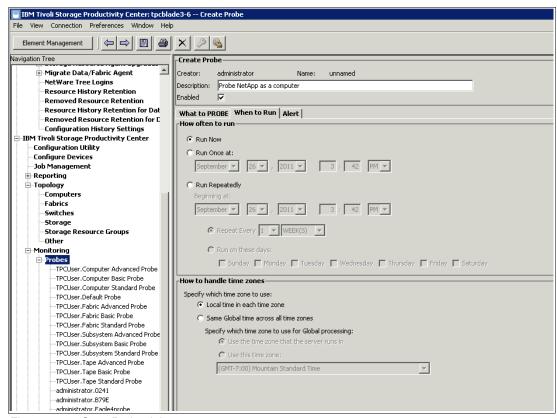


Figure 8-54 Save Probe Job

After a probe successfully completes, you can verify that the Tivoli Storage Productivity Center has the filer data by viewing the Tivoli Storage Productivity Center dashboard.

On the Monitored Server Summary panel, you see the total number of Network Appliance devices that are monitored and the total file system and disk capacities (Figure 8-55).

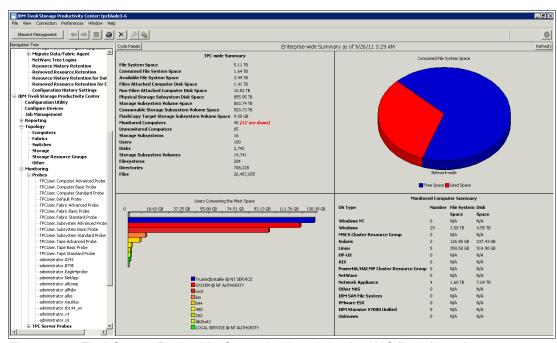


Figure 8-55 Tivoli Storage Productivity Center dashboard showing NAS filer information

Running a Scan job

To collect more detail information of file systems, files, and directories, you can run a Scan job. In the Tivoli Storage Productivity Center GUI Navigation Tree, expand **Data Manager** → **Monitoring**, right-click **Scans**, and click **Create Scan**. In the **Filesystems** tab, remove all other entries from the **Current Selections**, and add only the NAS filer to it, as shown in Figure 8-56.

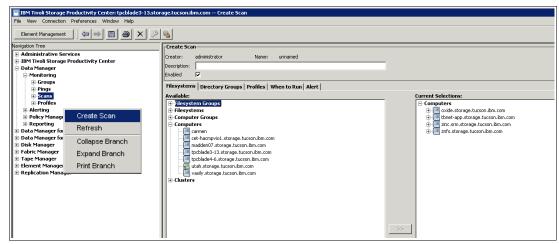


Figure 8-56 Define Scan job

In the **Profiles** tab, select all the default profiles and apply them to file systems and directories by clicking the double right angle bracket (>>) button (Figure 8-57). Profiles allow us to specify what statistical information is gathered and to fine-tune and control what files are scanned.

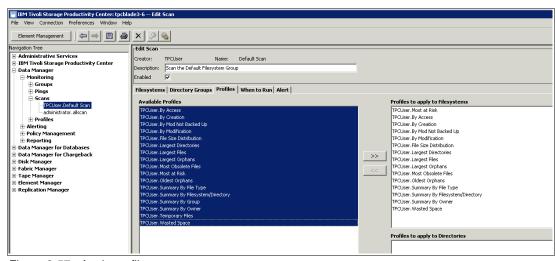


Figure 8-57 Apply profiles

Click **When to Run** tab, select the **Run Now** button, and click **Save** to save the scan job (Figure 8-58). Tivoli Storage Productivity Center asks for a name for the Scan job. Give a job name and click **OK** to start the scan job. You can right-click the job name and select **Update Job Status** to check the running job status. Wait until the scan job finishes.

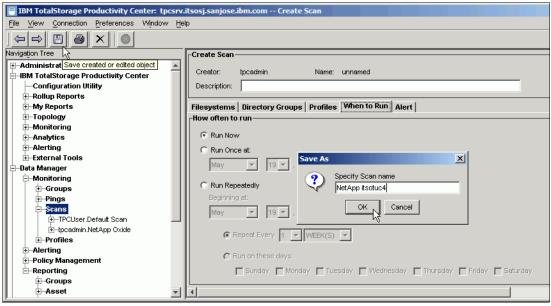


Figure 8-58 Save Scan job

8.6.3 Configuring the UNIX Storage Resource Agent

In this section, we document the procedure to configure Tivoli Storage Productivity Center to monitor the IBM N series or NetApp filer though a Windows server through a UNIX proxy Storage Resource Agent.

Checking NFS config on NAS filer

In the NAS FilerView, select NFS \rightarrow Manage Exports to check the NFS configuration. When NFS license is enabled, the root file system /vol/vol0 is exported. See Figure 8-59.



Figure 8-59 NAS NFS exports

Installing the Storage Resource Agent on a UNIX host

Install a Tivoli Storage Productivity Center Storage Resource Agent on the UNIX server. We use it as a proxy agent to collect information from the NAS.

See the 8.3, "Storage Resource Agent installation methods" on page 267 for additional details.

Mounting the NAS root file-system to a UNIX host

The root file system from the NetApp filer must be mounted on the agent where the Storage Resource Agent will be deployed. This way ensures that during post-installation discovery, the NetApp filer will be discovered automatically (see Figure 8-60).

Figure 8-60 Mount root file system from NAS filer

Discovering NAS devices in Tivoli Storage Productivity Center

Depending on your environment, the two ways to have NAS devices managed are as follows:

- ▶ Run a discovery to discover the NAS devices in your environment.
- ► Manually add the devices to Tivoli Storage Productivity Center.

We describe both methods.

Managing NAS devices through discovery

Complete the following steps:

- 1. Launch the IBM Tivoli Storage Productivity Center GUI.
- Expand the Administrative Services in the Navigation Tree and select Administrative Services → Discovery.

To run an NAS/NetApp Discovery job, right-click **Windows Domain**, **NAS and SAN FS**, and select **Run Now** (Figure 8-44 on page 304).

3. After the discovery job completes, but before IBM Tivoli Storage Productivity Center can perform operations (Probe/Scan) against NetWare, NAS, or NetApp filers, they must be licensed first. The reason is because in very large environments you might not want to automatically license all the discovered NAS, so you have a choice as to which servers to license.

Select the **Administrative Services** \rightarrow **Configuration** \rightarrow **License Keys**, and double-click the **TPC for Data** entry (Figure 8-45 on page 305).

The Licensing tab opens. Locate the NetWare, NAS, or NetApp filer you want. In the Licensed column, select the associated check box for the NAS filer, and then save changes by selecting the Disk icon on the toolbar (Figure 8-61).

Important: You must save the filer after you license the NAS filer before you leave this panel, otherwise Tivoli Storage Productivity Center does not save the change to its repository.

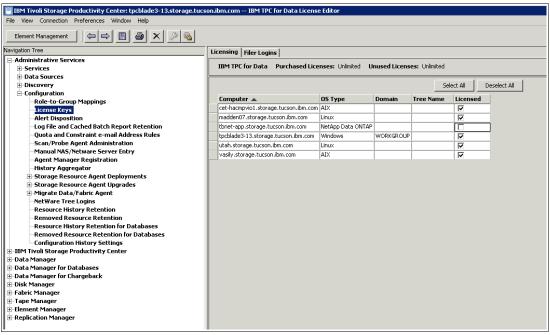


Figure 8-61 NAS Licensing tab

You do not need to use the **Filer Login** tab, which is required only for the Windows environment.

4. Run a discovery job again to get further information about the NAS filer. See Figure 8-44 on page 304.

Manually adding NAS device to Tivoli Storage Productivity Center

Complete the following steps:

- 1. Launch the Tivoli Storage Productivity Center GUI.
- Expand the Administrative Services in the Navigation Tree and select Administrative Services → Configuration → Manual NAS/Netware Server entry. The panel shown in Figure 8-62 opens.

Click **Add NAS Server** button, enter the following information, and click **OK**:

- Network name:
 - Enter the fully qualified host name or IP address of the NAS filer.
- Data Manager Agent OS Type:
 - Select the operating system of the computer that contains the agent that will gather information about the NAS filer. In our case, we select **Unix** here.
- Accessible from:
 - Select the UNIX agent host from drop-down list that you want to 'discover' the NAS filer.

– SNMP Community:

Enter the SNMP community name, the default is PUBLIC. This is used to get information from the NAS filer.

Login ID and Password:

Windows only. It is disabled when you click **Unix** as the Data Manager Agent OS Type.

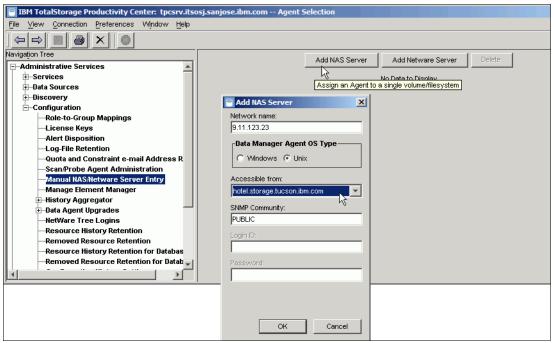


Figure 8-62 Manual add NAS server

If all the information you provided is correct, the NAS filer is added to the panel shown as Figure 8-63.

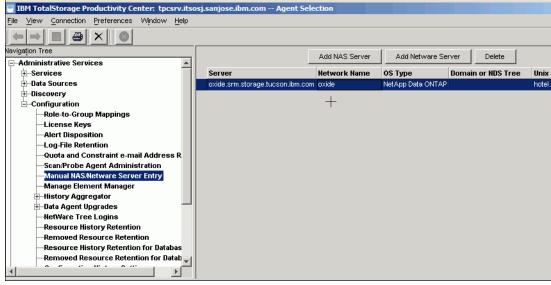


Figure 8-63 Manually added NAS filer

Setting the Scan/Probe Agent

After you manually added or discovered the NAS device, you need to set the Scan/Probe Agent before scanning or probing the filer. Because there is no Storage Resource Agent installed on the NAS filer itself, we need another agent running the Scan/Probe job as a proxy agent instead. After this task is done, Tivoli Storage Productivity Center will treat these devices as normal servers with attached storage. An agent can scan multiple NAS Servers, and NAS Server can be scanned by multiple agents, so we can set up parallel scans.

Remember: With this step you are assigning workload to computers. These scans create traffic on the IP network. We suggest that you select proxy agents on the same network as the NAS filer. This will maximize performance and minimize network load on other parts of your infrastructure. Scans over the network are not as performant as local scans. The scan might take extra time to run if it is doing a proxy scan on a large NAS device. For large NAS devices, use multiple proxy agents. Normally the default scan of once per day is not required, therefore, consider weekly scans.

Select Administrative Services \rightarrow Configuration \rightarrow Scan/Probe Agent Administration. Click the NAS filer that you want to define the Scan/Probe agent. Click the NAS filer entries, and click **Set agent per row**. In the pop-up window, choose the UNIX Storage Resource Agent that has the NAS filer attached (see Figure 8-64).

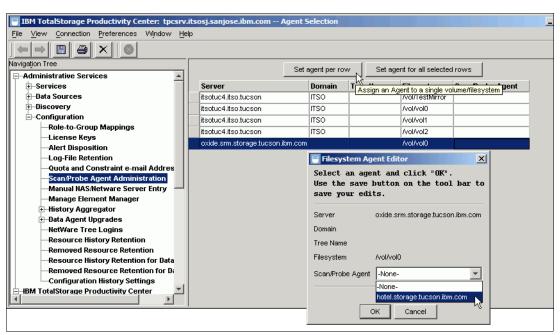


Figure 8-64 Set Scan/Probe agent

Now, you see that the NAS filer file systems have been assigned the Scan/Probe Agents. Make sure you click **Save** in the toolbar to save the change (Figure 8-65).

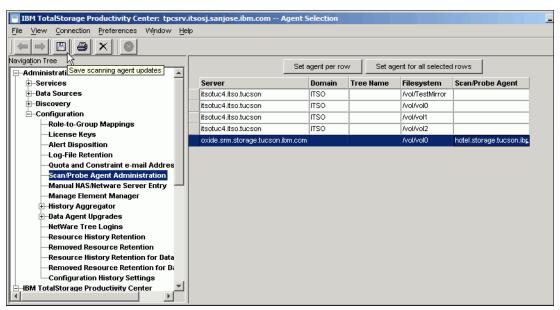


Figure 8-65 Save Scan/Probe agent

Running a Probe job

After the Scan/Probe Agent has been set, a probe to collect the devices hardware configuration and file system information are run automatically. If you want to create and run an additional probe, you can so do. In the Tivoli Storage Productivity Center GUI Navigation Tree, select **IBM Tivoli Storage Productivity Center** \rightarrow **Monitoring** \rightarrow **Probes**. Right-click and select **Create Probe**. Click **What to PROBE** tab, remove all other entries from the Current Selections, and add the NAS filer only to it, as shown in Figure 8-66.

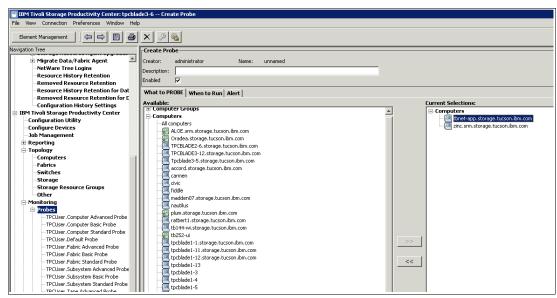


Figure 8-66 Define a Probe job

Click **When to Run** tab, and choose the radio button of **Run Now**, and save the probe job by clicking **Save** on the toolbar (see Figure 8-54 on page 310), the probe jobs starts. You can right-click the Probe Job and select **Update Job Status** to check the running job status. Wait until the probe job finishes.

After a probe successfully completes, you can verify that Tivoli Storage Productivity Center has the filer data by viewing the Tivoli Storage Productivity Center dashboard, shown in Figure 8-67.

The Monitored Server Summary Panel shows the total number of Network Appliance devices monitored and the total file system and disk capacities (Figure 8-67).

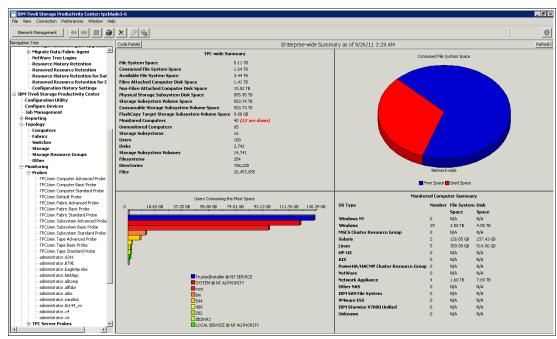


Figure 8-67 Tivoli Storage Productivity Center Dashboard showing NAS Filer Information

Running a scan job

To collect more detailed information of file systems, files, and directories, you can run a scan job. In the Tivoli Storage Productivity Center GUI Navigation Tree, expand **Data Manager** → **Monitoring**, right-click **Scans**, and click **Create Scan**. In the Filesystems tab, remove all other entries from the **Current Selections**, and add only the NAS filer to it, as shown in Figure 8-68.

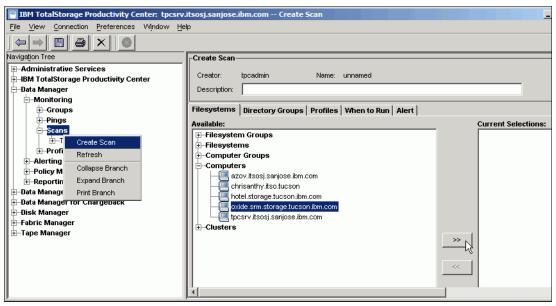


Figure 8-68 Create Scan job

In the **Profiles** tab, we select all of the default profiles and apply them to file systems and directories by clicking the >> button, see Figure 8-57. Profiles allow us to specify what statistical information is gathered and to fine-tune and control what files are scanned.

Click **When to Run** tab, select the **Run Now** button, and save the scan job by clicking the Save button on the toolbar (see Figure 8-69). Tivoli Storage Productivity Center asks for a name for the scan job. Enter a job name and click **OK** to start the scan job. You can right-click the job name and select **Update Job Status** to check the running job status. Wait until the scan job finishes.

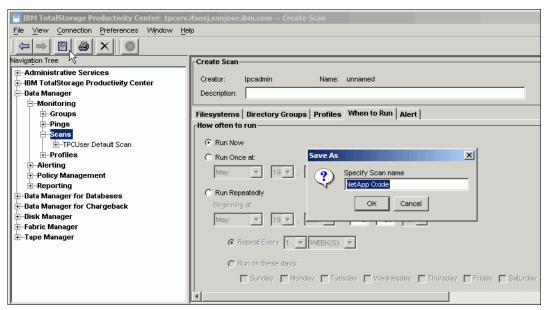


Figure 8-69 Save scan job

8.6.4 Retrieving and displaying data about the NAS filer

We have now set up the basic NAS collection jobs, we can start to view the information about NAS filers. Actually after setting Scan/Probe Agent for NAS filers, Tivoli Storage Productivity Center will treat these devices as normal computers with attached storage. For detailed information, see Chapter 8, Getting "Started with Tivoli Storage Productivity Center," sections 8.6 and 8.7 in *IBM Tivoli Storage Productivity Center: The Next Generation*, SG24-7194.

In the following section, we show examples of how to retrieve and display data for NAS filer.

Viewing the NAS filer from the Topology View

Start viewing NAS filer information about the Topology Viewer by expanding **IBM Tivoli Storage Productivity Center** \rightarrow **Topology** \rightarrow **Computers**. Then click the plus sign (+) in the top-right corner of the Computers (unknown) box. You see the NAS filer is there (Figure 8-70). Click the NAS filer, from the tabular view, you can see its OS type is NetApp Data ONTAP.

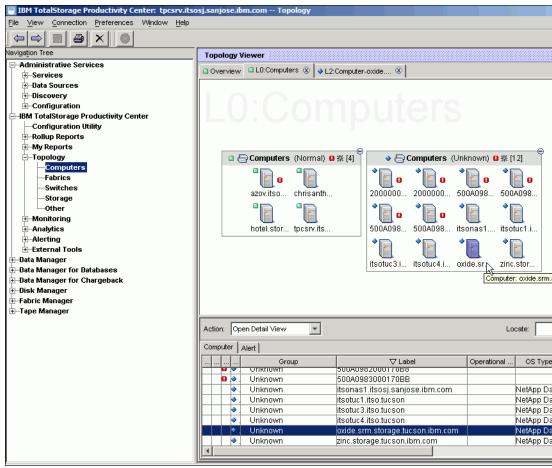


Figure 8-70 Topology view for computer

Double-click the NAS filer to view more details about this NAS filer from the L2: Computer view (Figure 8-71).

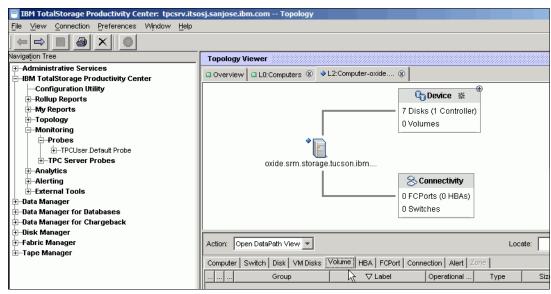


Figure 8-71 Topology view for NAS Filer

Navigation Tree-based asset reporting

Expand Data Manager \rightarrow Reporting \rightarrow Asset \rightarrow By OS Type \rightarrow Network Appliance. You can see the NAS filer asset information (Figure 8-72).

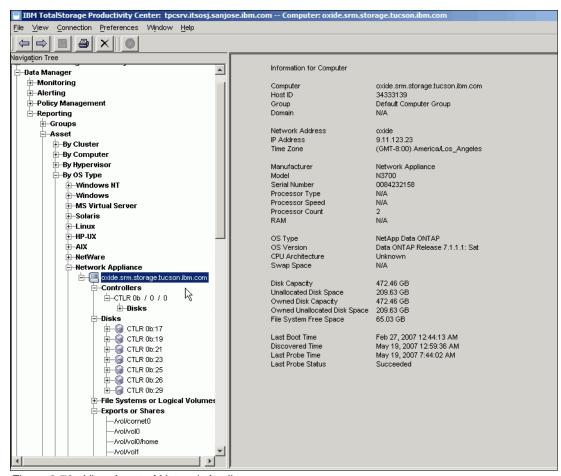


Figure 8-72 View Asset of Network Appliance

File system reporting

We can also generate reports from the NAS file systems. The following example shows how to generate a report:

 Expand Data Manager → Reporting → Capacity → Filesystem Free Space → By Computer. Click Generate Report as shown in Figure 8-73.

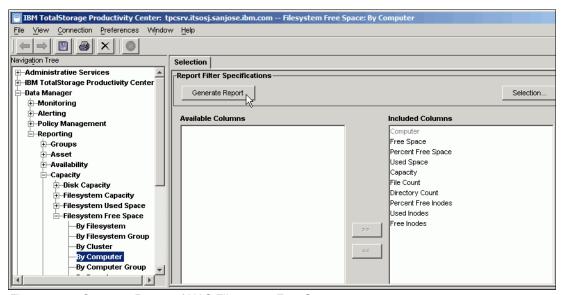


Figure 8-73 Generate Report of NAS Filesystem Free Space

2. In the next panel, select the NAS filer and click its magnifier icon (Figure 8-74).

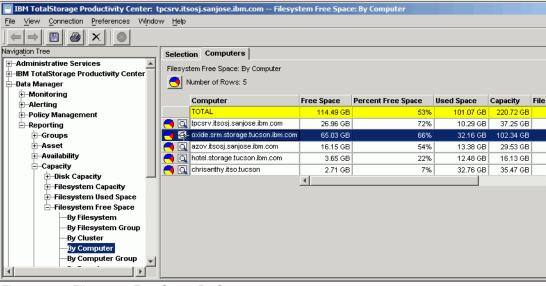


Figure 8-74 Filesystem Free Space By Computer

3. In the next panel, select the mount point you are interested in, right-click it, and then select **Chart space usage for selected** as shown in Figure 8-75.

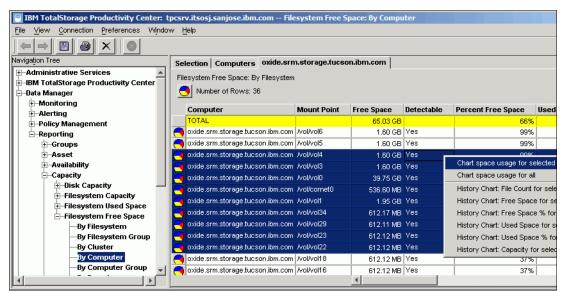


Figure 8-75 Filesystem Free Space from the NAS filer

The Filesystem Free Space chart is presented in the next panel. This chart shows the current free space on each volumes on the NAS filer. You can right-click the chart and click **Customize this chart** to customize the chart. On the pop-up panel, we select **4** in the **Maximum number of charts or series per screen** drop-down menu (Figure 8-76).

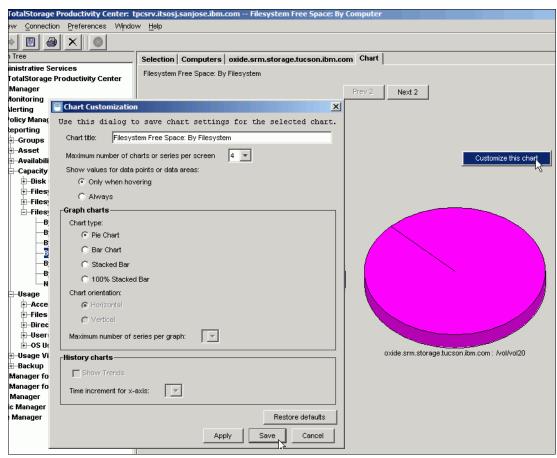


Figure 8-76 Chart of Filesystem Free Space by Computer

Now you can see the customized Chart of Filesystem Free Space By Computer as shown in Figure 8-77. You can click **Prev** and **Next** to see more charts of other NAS volumes.

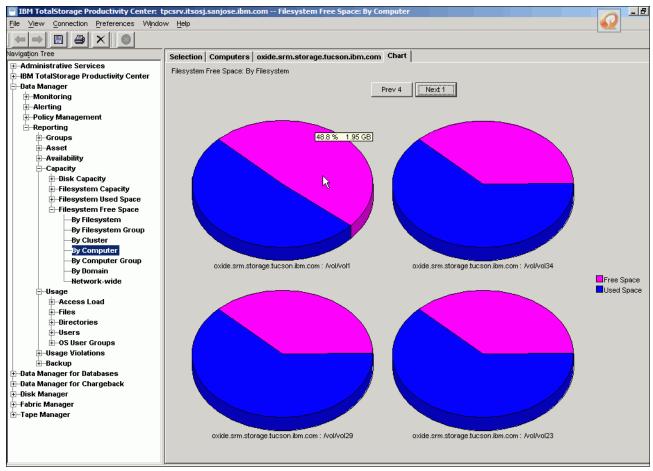


Figure 8-77 Customized Chart of Filesystem Free Space by Computer

NAS device quotas

You can import quotas that are set up on your NAS device into Tivoli Storage Productivity Center, as shown in Figure 8-78.

To work with these quotas, use the following tasks:

- 1. Run a scan job on the filers on which you want to import quotas.
- 2. Expand Data Manager \rightarrow Policy Management \rightarrow Network Appliance Quotas.
- 3. Right-click Schedules and select Create NetApp Quota Job.
- 4. Select a filer from which to import the quotas.
- 5. Under the Alert tab, you can define a condition that will trigger an alert if a certain percentage of the quota limit is reached.

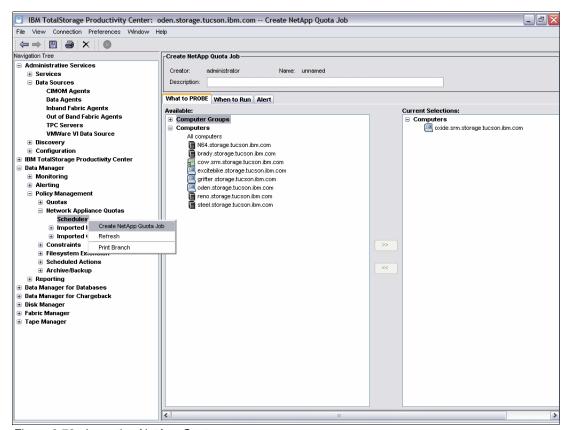


Figure 8-78 Importing NetApp Quotas

8.7 VMware support

The Storage Resource Agents can be installed on virtual machines on VMware ESX and ESXi 3.0.x, 3.5.x, 4.0.x, and 4.1.x.

Tivoli Storage Productivity Center Storage Resource Agents are able to gather virtual computer disk and file system information. To generate complete capacity reports, a Storage Resource Agent is required on each virtual machine. These agents are supported on the following system:

- ► Windows 2003/2008
- ▶ Red Hat Enterprise Linux 4 and 5
- SUSE Linux 9, 10, and 11

8.8 VMware virtual machine reporting

To view detailed information about a particular virtual machine, select the virtual machine under **Data Manager** \rightarrow **Reporting** \rightarrow **Asset** \rightarrow **By Computer**.

Attention: To see detailed information about the virtual machines, you must have a Storage Resource Agent deployed on each virtual machine.

The report in Figure 8-79 shows detailed information regarding the machine's assets:

- ▶ Machine Hostname
- Host ID: Unique machine identifier generated by the Tivoli GUID
- Group and Domain Information
- Network Address, IP Address
- ► Machine Time Zone
- Manufacturer, Model and Serial Number
- ► Processor Type, Speed and Count
- ► RAM Information
- Operating System Type and Version
- ► CPU Architecture and Swap Space
- ► Disk Capacity, Unallocated Disk Space
- ► Filesystem Free Space
- ► Last Boot Time, Discovered Time
- ► Last Probe Time and Status
- ► For VMWare virtual machines, information regarding hypervisor and VM configuration file.

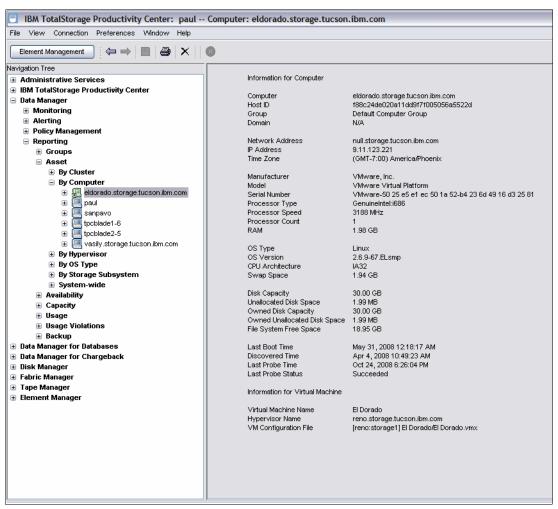


Figure 8-79 Virtual Machine Asset report

From this view, you can drill down into particular virtual machines, controllers, disks, file systems, exports and shares, and monitored directories.

To view details regarding disks assigned to a virtual machine select **Data Manager** \rightarrow **Reporting** \rightarrow **Asset** \rightarrow **By Computer** \rightarrow **[Computer Name]** \rightarrow **Disks** \rightarrow **[Disk #]**. The disk detail panel contains four tabs: General, Paths, Latest Probe, and Probe History:

- The General page (Figure 8-80) includes the computer name, path name, SCSI target ID, logical unit number and the number of access paths. This page also includes disk information such as the manufacturer, model number, firmware, serial number and manufacture date of the disk.
- The Paths page shows information regarding the host, OS type, path, controller, instance, bus number, SCSI target ID, and logical unit number.
- ► The Latest Probe page shows information gathered by Tivoli Storage Productivity Center during the most recent probe of the disk. This page includes information about the sectors, number of heads, number of cylinders, logical block size, disk capacity, RPM information, power-on time, failure prediction, disk defect information and time of last probe.
- ► The Probe History page shows the history of probes that have been run on this disk for tracking purposes.

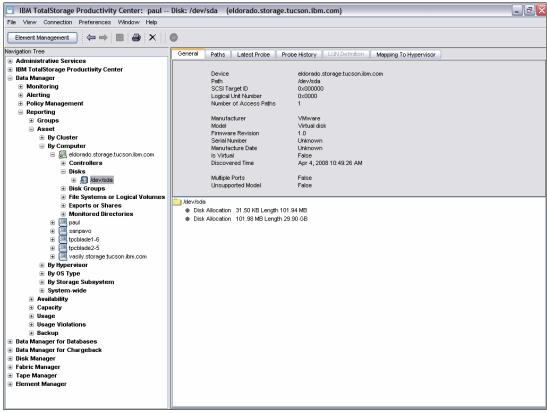


Figure 8-80 Virtual Machine Disk Information

For a given virtual machine disk, you can also view how it is mapped from the hypervisor. To do so, select the Mapping to Hypervisor tab on the disk information report.

8.9 Batch reporting

The batch reporting feature enables you to run any report on a regularly scheduled basis. This feature enables you to conveniently run reports and gather data on a set schedule.

To create the batch report, perform the following steps:

- 1. Select IBM Tivoli Storage Productivity Center → My Reports → Batch Reports.
- 2. Right-click and select Create Batch Report, as shown in Figure 8-81.

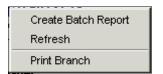


Figure 8-81 Batch report options

- 3. Select the report type, as shown in Figure 8-82. You can select the following reports:
 - Asset System-wide
 - Storage Subsystems
 - Availability
 - Capacity
 - Usage
 - Usage Violations
 - Backup
 - Groups

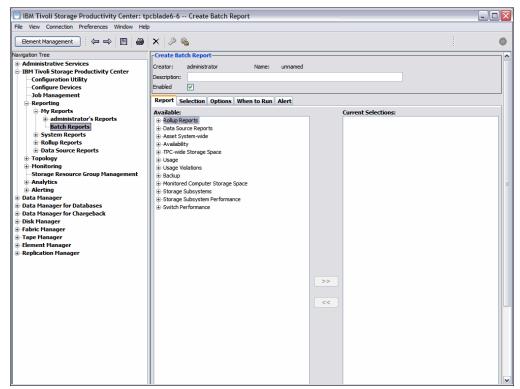


Figure 8-82 Create Batch Report

Based on the report type, the selections panel gives you the ability to narrow down the columns in the report or filter the report to include specific entities (Figure 8-83).

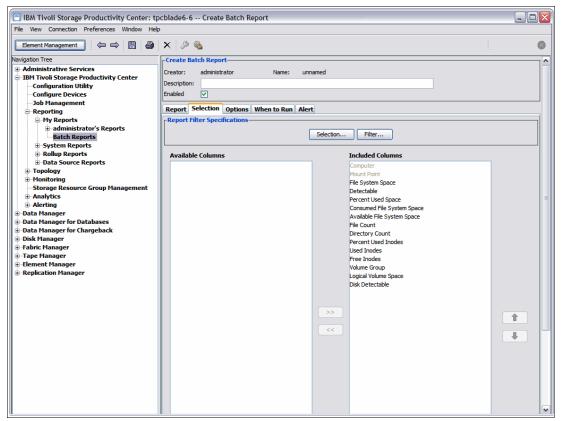


Figure 8-83 Batch Report Selection

The Options panel has been enhanced in Tivoli Storage Productivity Center V4.2. You can now either generate the batch reports on the Tivoli Storage Productivity Center server machine or specify a custom location for the generated batch reports. You can also select the type and format of report to generate in this panel (see Figure 8-84).

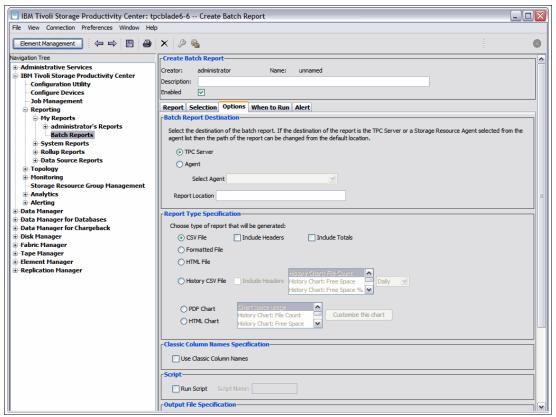


Figure 8-84 Batch Report options

On this panel, you can also specify the format for the output file name (Figure 8-85).

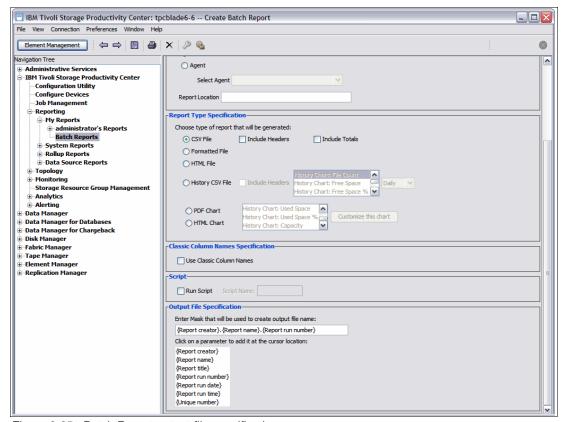


Figure 8-85 Batch Report output file specifications

In the next panel (Figure 8-86) you can specify when to run the batch report, and set up a schedule to run the defined batch report repeatedly.

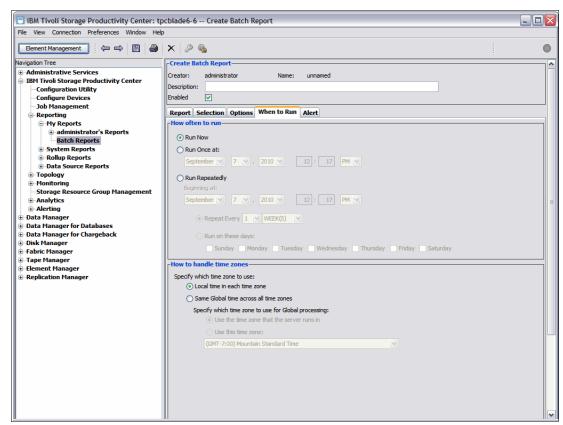


Figure 8-86 Batch Report When to Run panel

In the Alert panel, you can specify an alert to be generated if the batch report generation fails (see Figure 8-87).

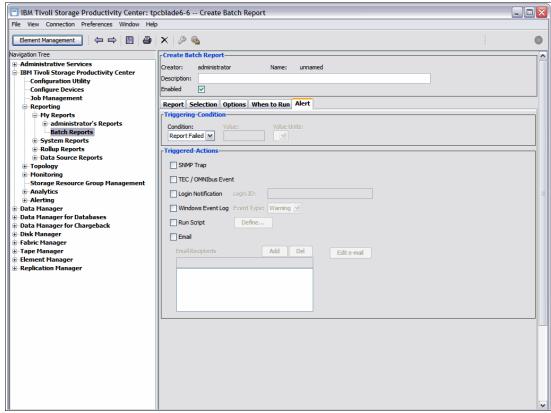


Figure 8-87 Batch Report alert definition

8.10 Storage Resource Agent fabric function

With Tivoli Storage Productivity Center V4.2, certain fabric functions are incorporated into the Storage Resource Agent:

- Collection of location HBA information
- ► Fabric discovery, topology, and zoning probes through inband GS-3 commands
- Fabric zone control for Cisco, QLogic, and McDATA fabrics through inband GS-3 commands
- New Agent Assignment feature

Support: By changing zone configuration and reporting of HBA, fabric topology and zoning information is not supported on hosts running Linux on IBM pSeries® or IBM zSeries®. These functions are also not supported on VMware guest images or SLES 11 on x86-64 with Emulex HBAs.

8.10.1 HBA library requirements

Refer to the Tivoli Storage Productivity support matrix to see the HBA driver and firmware library requirements.

For a detailed list of supported platforms, see the Tivoli Storage Productivity Center support site:

http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition

Click the **Documentation** link, and enter Platform Support: Agents, Servers and GUI in the Search support box. Click the link to the Tivoli Storage Productivity Center 4.2x Supported Storage Matrix.

8.10.2 Storage Resource Agent fabric enhancements

The Storage Resource Agents have been enhanced to support the following functionalities:

- ▶ Distinct fabric discovery and fabric probe functions in Storage Resource Agent. Fabric discovery processing is part of computer probe processing
- ► Separate Storage Resource Agent fabric probe jobs to collect topology and zoning information to improve granularity
- ► Storage Resource Agents automatically assigned to probe specific fabrics based on requirements. Fabric functionality can be enabled or disabled for a particular Storage Resource Agent in the data sources panel
- Automatic zone probe is not performed after zone changes
- More targeted probes of fabrics (checking in fabric is designated as a fabric to be probed by user, determining fabrics that need to be probed for events, agents to be used for probe)
- Support for Brocade HBAs (Support for multiple-port HBAs using CTPassThru Version 2 commands)
- ► Retry logic at various points in flow for robustness (retries on GS-3 commands, using V1 and V2 versions of CTPassThru, retries in communication with Storage Resource Agent, and so on)

8.10.3 Fabric agent assignment

Tivoli Storage Productivity Center V4.2 now has an intelligent *Agent Assignment* algorithm to identify the minimum set of agents that are required to gather data for a fabric. Therefore, you no longer have to be concerned about causing *fabric storms* by having too many agents on a single fabric. This way allows for failover to alternate agents during fabric probe and CIM indications to improve resiliency.

8.11 Agent resource utilization

The Storage Resource Agent in Tivoli Storage Productivity Center V4.2 has significant performance improvements compared to the legacy data agents available in previous releases.

The time required to deploy a Storage Resource Agent is significantly less than the data agent installation time. On certain platforms, the deployment time has been reduced by 500%.

The average probe and scan times for a Storage Resource Agent is comparable to those of a Data agent on the same hardware. However, the memory consumption for the Storage Resource Agent is significantly less than the legacy agents. A typical Storage Resource Agent scan process consumes 5 MB on Windows 2003, compared to 50 MB with the legacy agent.

The Storage Resource Agent consumes minimal CPU and memory resources at deployment time and when it is running idle on the servers (in daemon mode). For this reason, use the daemon Storage Resource Agents.

8.12 HBA information reports

To view information regarding the HBAs installed in your Storage Resource Agents, select **Administrative Services** \rightarrow **Data Sources** \rightarrow **Data/Storage Resource Agents**, and click the magnifying glass of a particular agent (Figure 8-88).

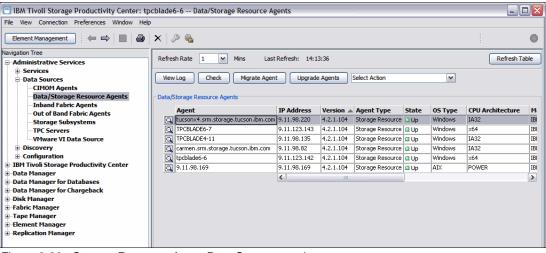


Figure 8-88 Storage Resource Agent Data Source panel

Click the magnifying glass. An agent detail panel opens. It lists the properties of the installed HBAs (Figure 8-89).

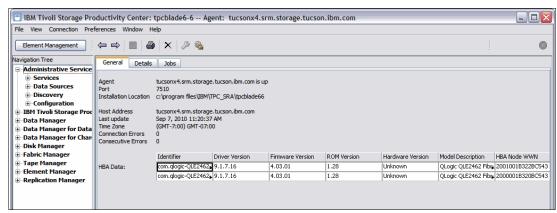


Figure 8-89 Storage Resource Agent HBA properties

You can also run a DB2 query to generate an output of all of the HBAs installed in your environment. See Example 8-1 for details about the query. The report can be generated using either the DB2 command-line interface or the command center.

Tip: You can also use a reporting tool, such as BIRT, to generate this report.

Example 8-1 DB2 query for HBA report

```
select
  CS.DISPLAY NAME,
  CS.IP ADDRESS,
  HBA.NAME,
  HBA.SERIAL NUMBER,
  HBA.INSTANCE NUMBER,
  HBA.DRIVER VERSION,
  HBA.FIRMWARE VERSION,
  HBA.ROM VERSION,
  HBA.HW VERSION,
  HBA.MODEL,
  HBA.WWNN,
  HBA.WWPN,
  HBA.BUS NUMBER
from TPCREPORT.COMPUTERSYSTEM CS right join TPCREPORT.HOSTBUSADAPTER HBA on
CS.COMPUTER_ID = HBA.COMPUTER_ID and HBA.DETECTABLE='True'
where CS.DETECTABLE='True'
```

8.13 Collecting Storage Resource Agent support data

With Tivoli Storage Productivity Center V4.2, you can now collect the Storage Resource Agent support information through a remote job on the Tivoli Storage Productivity Center GUI, without having to log in to the individual Storage Resource Agent computers.

To collect the Storage Resource Agent support information, click **Administrative Services** → **Data Sources** → **Data/Storage Resource Agents**, select the Storage Resource Agent that you want to collect the data from, and select **Collect Service Data** on the drop-down selection (Figure 8-90).

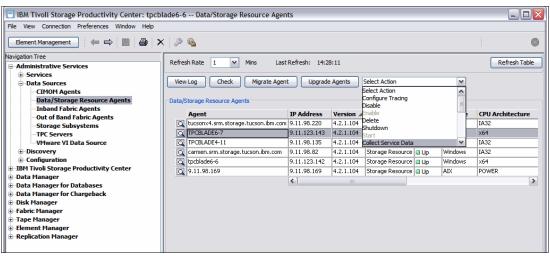


Figure 8-90 Storage Resource Agent service data collection

This step invokes the service collection script on the Storage Resource Agent, compresses the contents, and sends it to the Tivoli Storage Productivity Center server. Click **OK** at the confirmation (Figure 8-91).



Figure 8-91 Storage Resource Agent Service Data Collection confirmation

8.14 Clustering support

Storage Resources agents can be installed on Microsoft Cluster Server (MSCS) and IBM High Availability Cluster Multiprocessing (IBM HACMP™) computers.

Support: Clustered scan support is limited to daemon Storage Resource Agents.

Legacy data agents installed on cluster nodes should be migrated to daemon Storage Resource Agents.



Tivoli Storage Productivity Center Select and Tivoli Storage Productivity Center for Disk Select

In this chapter, we provide an overview of Tivoli Storage Productivity Center Select and Tivoli Storage Productivity Center for Disk Select. We compare them to Tivoli Storage Productivity Center Standard Edition and to Tivoli Storage Productivity Center for Disk Midrange Edition.

9.1 Overview

Tivoli Storage Productivity Center for Disk Midrange Edition was introduced in June, 2010 with the same functionality as the Tivoli Storage Productivity Center for Disk, but with a different licensing type, targeting entry-level and midrange storage environments.

Starting with Version 4.2.2, this new licensing type was extended to the Tivoli Storage Productivity Center family to also cover the other functions provided by Tivoli Storage Productivity Center Standard Edition. This new license was named Tivoli Storage Productivity Center Select. At the same time, Tivoli Storage Productivity Center for Disk Midrange Edition was renamed to Tivoli Storage Productivity Center for Disk Select.

Tivoli Storage Productivity Center Select and Tivoli Storage Productivity Center for Disk Select provide the same features and functions as Tivoli Storage Productivity Center Standard Edition and Tivoli Storage Productivity Center for Disk respectively, but are limited to managing IBM XIV, IBM SAN Volume Controller, IBM Storwize V7000, and IBM System Storage DS3000, DS4000, and DS5000 devices. From the software installation and utilization perspective, there are no differences in the processes compared to non-Select products.

9.2 Supported devices and firmware levels

Tivoli Storage Productivity Center Select and Tivoli Storage Productivity Center for Disk Select support the IBM System Storage Devices including XIV, SAN Volume Controller, Storwize V7000, DS3000, DS4000, and DS5000.

Support: SAN Volume Controller is supported, but only if the back-end storage is from IBM XIV, Storwize V7000, DS3000, DS4000 or DS5000 devices. IBM SAN Volume Controller boxes are not counted as Storage Devices when computing Tivoli Storage Productivity Center pricing.

View the supported firmware levels at the Tivoli Storage Productivity Center support site:

http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition

Click the **Documentation** link, and enter Supported Storage Products Matrix in the Search support box. Click the link to the Tivoli Storage Productivity Center 4.2 document.

9.3 Licensing methodology

Tivoli Storage Productivity Center Select and Tivoli Storage Productivity Center for Disk Select offer highly competitive pricing for the devices and a pricing metric based on the number of *storage devices* or *enclosures*. The license price is based on the number of managed enclosures, not their respective capacities.

A Storage Device, for licensing purposes with Tivoli Storage Productivity Center Select and Tivoli Storage Productivity Center for Disk Select, is defined as an independently powered, channel attached device that stores or controls the storage of data on magnetic disks or solid state drives, such as disk controllers and their respective expansion units, each constituting separate Storage Devices. In this case, each supported controller enclosure, even if they contain no disk drives, and each expansion enclosure to be managed by this licensing type, counts as one Storage Device.

Tivoli Storage Productivity Center for Disk Select performance monitoring is supported on XIV, Storwize V7000, DS3000, DS4000, and DS5000 Storage Devices and includes Basic Edition capabilities. Tivoli Storage Productivity Center Select includes Standard Edition capabilities for the same IBM storage devices. SAN Volume Controller performance monitoring is supported with the Select licenses when the back-end storage is made up of the supported IBM storage devices.

Tivoli Storage Productivity Center Select and IBM Tivoli Storage Productivity Center for Disk Select are set apart from Tivoli Storage Productivity Center Standard Edition and IBM Tivoli Storage Productivity Center for Disk respectively in the following ways:

- ► They are designed to support IBM entry-level and midrange storage disk products (IBM XIV, Storwize V7000, DS3000, DS4000 and DS5000 and IBM SAN Volume Controller virtualizing these products).
- ► They are licensed per storage device, such as disk controller enclosures and their respective expansion units. IBM System Storage SAN Volume Controller, IBM System Storage SAN Volume Controller Entry Edition, IBM Virtual Disk System controllers, or both, are not counted as a storage device when computing Select pricing within these environments.

9.4 Key benefits

Tivoli Storage Productivity Center Select and Tivoli Storage Productivity Center for Disk Select support IBM entry-level and midrange storage disk products (IBM XIV, Storwize V7000, DS3000, DS4000 and DS5000). They are licensed per storage device and not per terabytes, which makes them in most cases more suitable for midrange environments when pursuing Storage Management capabilities at more affordable prices.

Key benefits of Tivoli Storage Productivity Center for Disk Select are as follows:

- ▶ Provides reporting across multiple midrange arrays from a single console.
- ► Helps monitor metrics such as throughput, input and output (I/O) rates, data rates, and cache utilization.
- Receives timely alerts that can enable event action based on your policies when thresholds are exceeded.
- Offers continuous and proactive performance analysis with comprehensive real-time monitoring and fault identification to help improve SAN availability.
- ► Helps you improve storage return on investment (ROI) by helping to keep SANs reliably and dependably operational.
- Helps reduce storage administration costs by simplifying the management of complex SANs.
- Supports the performance reporting capabilities on the IBM System Storage SAN Volume Controller (SVC) and SVC Entry Edition with attached XIV, Storwize V7000, DS3000, DS4000 and DS5000 devices.
- Supports performance reporting capabilities for any storage virtualized by the IBM Storwize V7000.

Tivoli Storage Productivity Center Select includes all the foregoing benefits for Disk Management, as well as these:

- Data Management functions, same as provided by Tivoli Storage Productivity Center for Data.
- ► Advanced analytics for storage provisioning, change management and performance reporting capabilities.
- ► Additional management, control and performance reporting for the SAN infrastructure.
- Automated device discovery, topology rendering, error detection and fault isolation, SAN error predictor, zone control, monitoring and alerts, and event management for heterogeneous enterprise SAN environments.
- ► Integrate with Tivoli Productivity Center for Replication to provide a single control for both replication and storage functions.
- ► Automated end to end provisioning capability that enables administrators to provision storage in a wizard driven approach with guidelines based on the storage system past performance.
- ► Help simplify the management and improve the availability of the SAN environment.
- ► Automatic device discovery and allow multiple SAN views, including physical, logical and zone views.
- ► View and analyze multiple aspects of the storage environment, including capacity, utilization, assets and availability.
- Detect storage events and generate the appropriate alerts to the administrator.
- Basic diagnostic capabilities to show which resources are impacted by an availability or performance issue in the SAN.



Tivoli Storage Productivity Center for Replication

In this chapter, we describe the new features, functions, and enhancements of Tivoli Storage Productivity Center for Replication V4.2. We also highlight the changes since Tivoli Storage Productivity Center for Replication V4.1.1 and focus on the following changes:

- Tivoli Storage Productivity Center for Replication Open HyperSwap
- ► SAN Volume Controller and Storwize V7000 enhancements
- ► DS8000 enhancements
- ► XIV enhancements

10.1 New V4.2 features and functions

Tivoli Storage Productivity Center for Replication V4.2 introduces new features, functions, and enhancements:

- New Tivoli Storage Productivity Center for Replication functions:
 - Open HyperSwap replication
 - Copy set soft removal of hardware relationship
 - Log packages download from Tivoli Storage Productivity Center for Replication GUI
 - Path Manager
- ► SAN Volume Controller session enhancements:
 - Space-efficient volumes support
 - Incremental FlashCopy support
- ▶ DS8000 session enhancements:
 - Extent space-efficient volumes support
 - Global Mirror session enhancements
 - Multiple Global Mirror sessions support in a storage system
- ► XIV session enhancements:
 - Snapshot session support
 - Metro Mirror Failover/Failback session support
 - Global Mirror Failover/Failback session support

With V4.2, Tivoli Storage Productivity Center for Replication no longer supports DB2 as the data store for its operational data. Tivoli Storage Productivity Center for Replication uses an embedded repository for its operational data.

The Tivoli Storage Productivity Center for Replication V4.2 installation program can automatically migrate any Tivoli Storage Productivity Center for Replication operational data, in an existing Tivoli Storage Productivity Center for Replication DB2 database. The migration is done to the Tivoli Storage Productivity Center for Replication embedded repository as part of upgrading to Tivoli Storage Productivity Center for Replication V4.2 from an earlier version. Details about the migration are described in Chapter 5, "Migrating Tivoli Storage Productivity Center base code to current level" on page 161.

10.2 Open HyperSwap replication

The Tivoli Storage Productivity Center for Replication 4.2 introduces new Open HyperSwap replication, which adds high availability support to the existing Metro Mirror session type. It can automatically fail over I/O from the primary logical devices to the secondary logical devices if a primary disk storage system failure occurs. This function can be done with minimal disruption to the applications that are using the logical devices.

In the following sections, we describe more details about the Open HyperSwap replication method and the required prerequisites to establish it.

Replication: Open HyperSwap is a new replication method for AIX operating systems; the existing HyperSwap functions support only IBM System z®. Both Open HyperSwap and HyperSwap provide high availability of data if a primary disk storage system failure occurs.

10.2.1 Description

Open HyperSwap replication applies to both planned and unplanned site switches. When a session has Open HyperSwap enabled, an I/O error on the primary site automatically causes the I/O to switch to the secondary site without any user interaction and with minimal application impact.

In addition, while Open HyperSwap is enabled, the Metro Mirror session supports disaster recovery. If a write is successful on the primary site but is unable to get replicated on the secondary site, Tivoli Storage Productivity Center for Replication suspends the entire set of data consistency checking, thus ensuring that a consistent copy of the data exists on the secondary site. If the system fails, this data might not be the latest data, but the data should be consistent and allow the user to manually switch host servers to the secondary site.

You can control Open HyperSwap from any system running Tivoli Storage Productivity Center for Replication (AIX, Windows, Linux, or IBM z/OS®). However, the volumes that are involved with Open HyperSwap must be attached to an AIX system. The AIX system is then connected to Tivoli Storage Productivity Center for Replication.

Figure 10-1 shows an overview of the Open HyperSwap function.

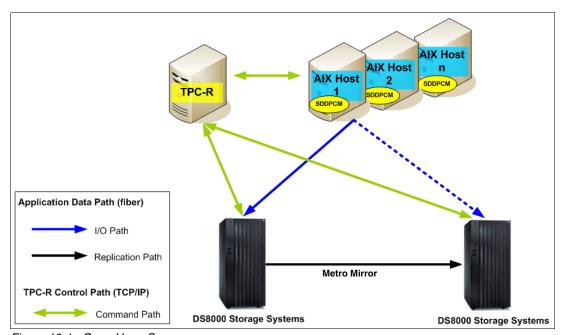


Figure 10-1 Open HyperSwap

10.2.2 Prerequisites

In this section, we discuss the prerequisites for Open HyperSwap.

Tivoli Storage Productivity Center for Replication requirements

If you want to use the Open HyperSwap session, you must have the Tivoli Storage Productivity Center for Replication Two Site Business Continuity license.

By default, Tivoli Storage Productivity Center for Replication uses TCP/IP port 9930 for communication with the AIX host for the Open HyperSwap and TCP/IP port 1750 for communication with DS8000 (HMC connection).

Details about other TCP/IP ports used by Tivoli Storage Productivity Center for Replication are described in *IBM Tivoli Storage Productivity Center Installation and Configuration Guide*, SC27-2337.

AIX requirements

Open HyperSwap support requires AIX Version 5.3 (with required APARs) or 6.1. You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at the following link:

http://www.ibm.com/support/docview.wss?&uid=swg21386446

You must have the following AIX modules installed:

- ► Subsystem Device Driver Path Control Module (SDDPCM) Version 3.0.0.0 or later
- ► Multi-Path Input/Output (MPIO) module (the version that is provided with AIX Version 5.3 or 6.1)

The TCP/IP connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server must be established.

SDDPCM details

SDD distinguishes the paths of the source volume from the paths of the target volume on an Open HyperSwap copy set. With an Open HyperSwap device, I/O can be sent only to the source volume, so when SDD selects paths for I/O, it selects only paths that are connected to the source volume. If no path on the source volume can be used, SDD initiates the Open HyperSwap request to Tivoli Storage Productivity Center for Replication and works together to perform the swap. After the swap, SDD selects the target volume paths for I/O.

AE daemon is new in SDDPCM and it is added to the SDDPCM installation package beginning with SDDPCM 3.0.0.0. The daemon is used for communication with Tivoli Storage Productivity Center for Replication to support Open HyperSwap.

SDDPCM host error log collection

If an unplanned situation or disaster recovery situation occurred, be sure to collect SDD host error data. Issue the **sddpcmgetdata** command, which will create a .tar file that includes all SDD host data, and is saved under the local directory when the command is issued on the host. Because the data in the log file is generated in real time, be sure to collect this data.

Important:

- Open HyperSwap does not yet support host clustering solutions such as Power HA (High Availability Cluster Multi-Processing, HACMP).
- Open HyperSwap device is not supported with a SAN boot volume group.
- Currently DB2 with raw device access is not supported.

You can find more details about SDDPCM in the *Multipath Subsystem Device Driver User's Guide*, GC52-1309.

Considerations:

- ► Tivoli Storage Productivity Center for Replication must not be installed on any AIX hosts that are involved in the Open HyperSwap session.
- ► For AIX 5.3, a single host can manage a maximum of 1024 devices when devices have been enabled for Open HyperSwap on the host, with 8 logical paths configured for each copy set in the session.
- ► For AIX 6.1, a single host can manage a maximum of 1024 devices when devices have been enabled for Open HyperSwap on the host, with 16 logical paths configured for each copy set in the session.

DS8000 hardware and license requirements

Open HyperSwap is supported only for IBM DS8000 storage systems. Version 5.1 or later microcode levels are supported.

You must also have Metro Mirror license on DS8000 storage systems because Open HyperSwap uses Metro Mirror replication.

Support: Open HyperSwap is supported only in Metro Mirror Failover/Failback replication for DS8000 storage systems.

10.2.3 Setting up Open HyperSwap session

Before using Open HyperSwap, you must set up your environment for this function. The general steps are as follows:

- 1. Prepare the AIX system for Open HyperSwap:
 - a. Install SDDPCM driver.
 - b. Check if the Arbitration Engine(AE) daemon is active.
 - c. Use the AIX configuration manager (**cfgmgr**) to identify all volumes that are involved with the Open HyperSwap session.
- 2. Set up the host connection of Tivoli Storage Productivity Center for Replication to the AIX system. Use the Tivoli Storage Productivity Center for Replication user interface to manually set up the connection to the AIX system. Use the Host Systems page to enter the IP address and port number for the AIX system.
- 3. Set up the Tivoli Storage Productivity Center for Replication Metro Mirror Failover/Failback session, selecting the Open HyperSwap function.
- 4. Add the copy sets to the session where all the volumes in the copy sets are volumes that are on the AIX system that is connected to Tivoli Storage Productivity Center for Replication.

Hosts: An AIX host can be part of only one Tivoli Storage Productivity Center for Replication session at a time. However, a Tivoli Storage Productivity Center for Replication session can manage multiple AIX hosts.

5. Start the Open HyperSwap session.

Scenario to set up Open HyperSwap

In the following scenario, we describe how to set up Open HyperSwap session:

1. After you install the SDDPCM driver, check the AE daemon by typing the command in Example 10-1. If the daemon is not active, start it by using the following start command:

startsrc -s AE/startAE

Example 10-1 SDDPCM AE daemon

jerome> lssrc	-s AE			
Subsystem	Group	PID	Status	
AE		12036	active	

2. Use the AIX configuration manager (cfgmgr) to identify all volumes that are involved with the Open HyperSwap session. You can check the volumes by issuing the SDDPCM command (Example 10-2). In our example, we use 20 volumes (10 primary and 10 secondary). Only two primary volumes are shown in Example 10-2.

Example 10-2 SDDPCM path query device command

3. To manage the Open HyperSwap, the AIX host must be connected to Tivoli Storage Productivity Center for Replication server. To add the AIX host, click **Host Systems** from the Tivoli Storage Productivity Center for Replication menu; the panel shown in Figure 10-2 opens.

Hosts and sessions: A single session that has Open HyperSwap enabled can manage multiple hosts. However, each host can be associated with only one session. Multiple hosts can share the same session.

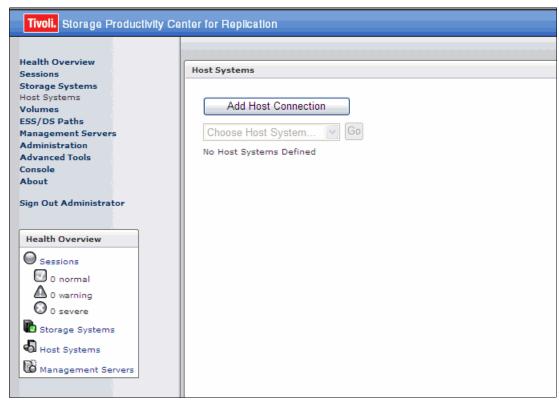


Figure 10-2 Add Host Connection panel

4. Click **Add Host Connection**. The next window opens (Figure 10-3). Enter the AIX host name or IP address and click **Add Host**.



Figure 10-3 Add Host Connection details

If the connection from Tivoli Storage Productivity Center for Replication and the AIX host is established, the AIX host is added to Tivoli Storage Productivity Center for Replication (Figure 10-4).

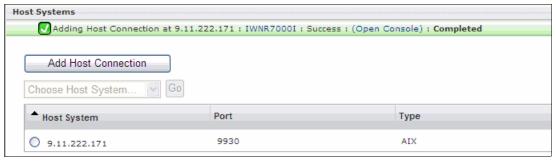


Figure 10-4 Host added to Tivoli Storage Productivity Center for Replication

 After you add the host connection, click Sessions (Figure 10-5). An overview of all defined sessions is listed. At this point, there are no defined sessions. Click Create Session to continue.



Figure 10-5 Create sessions

6. In the Create Session window (Figure 10-6), select **Metro Mirror Failover/Failback** and click **Next**. This is the only supported session with Open HyperSwap function.

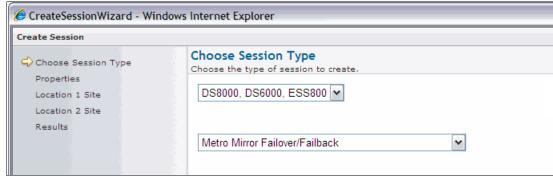


Figure 10-6 Define Metro Mirror Failover/Failback session

7. The Properties panel opens (Figure 10-7). The Properties panel is important because it gives you the option to specify this session with Open HyperSwap function. If you select this option, Open HyperSwap is triggered and redirects application I/O to the secondary volumes, when a failure occurs on the host-accessible volumes. You can also specify to prevent a HyperSwap from occurring by command or event by selecting Disable Open HyperSwap. In certain situations, you might have to temporarily disable Open HyperSwap capabilities for a session.

Within the Properties panel, you can also select one of the following options in the Metro Mirror Suspend Policy area:

- Hold I/O after Suspend:

Select this option if you want to block the application from writing while a consistent copy of the data is formed on the remote site. However, this option does not automatically release the application. This option keeps the source equal to the target. You must use the Release I/O command on the session or wait for the Hardware Freeze Timeout Timer to expire before the application can continue to write to the source.

Release I/O after Suspend:

Select this option if you want to block writes to the application while a consistent copy of the data is formed on the remote site, followed immediately by releasing the block so that the application can continue writing to the source. This option allows for little application impact, but causes the source to potentially differ from the target. This option is the default setting for all new sessions.

After you define the properties, click Next.

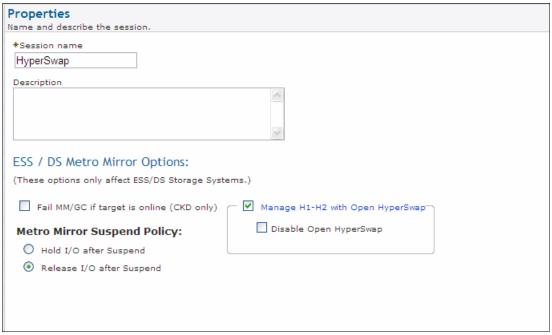


Figure 10-7 Metro Mirror Failover/Failback properties

8. In the panels shown in Figure 10-8 and Figure 10-9, you can define your site locations.

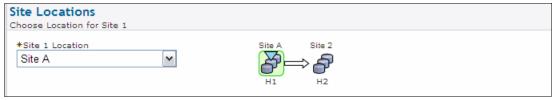


Figure 10-8 Define site location for Site 1

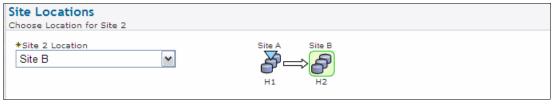


Figure 10-9 Define site location for Site 2

9. After you click **Next**, the session was successfully created (Figure 10-10).



Figure 10-10 Successfully created session

10. Click Finish and your session is listed in the Sessions overview panel (Figure 10-11).



Figure 10-11 Open HyperSwap session

11. After the Metro Mirror session with Open HyperSwap function is created, you have to populate this session with copy sets that consist of one H1 and one H2 Metro Mirror volumes. In the Sessions overview panel, select the session name radio button and select **Add Copy Sets** from the Select Action pull-down menu (Figure 10-12). Click **Go** to invoke the Add Copy Sets wizard.

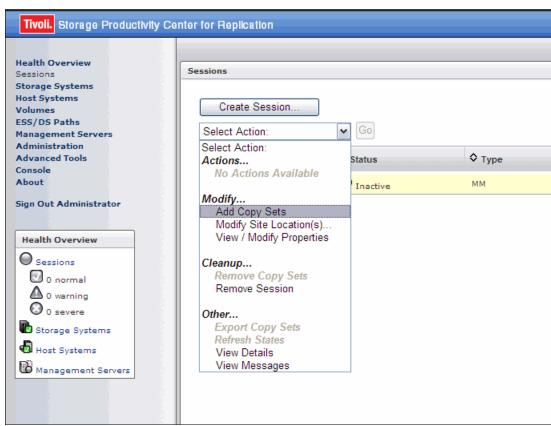


Figure 10-12 Add Copy Sets action

12. The Add Copy Sets wizard provides details about the primary volumes or local volumes, which are called Host 1 volumes, relating to the fact that these volumes reside in the Site 1. Select the desired Host 1 storage subsystem from the pull-down menu and wait for a few seconds to get the Host 1 logical storage subsystem list. Select the logical subsystem (LSS) where your H1 volume resides, and then choose the appropriate volumes from the Host 1 volume pull-down list.

The alternative way to add a large number of volumes to this session is to create a CSV file. If you have a CSV file ready, select **Use a CSV file to import copy sets** check box and provide a path to your CSV file.

In our example, we selected a DS8000 disk subsystem and all volumes from selected LSS as shown in Figure 10-13. Click **Next** to continue.

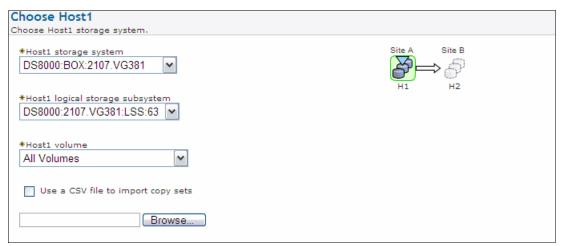


Figure 10-13 Add copy set to Metro Mirror session: Host1

13.In the next panel (Figure 10-14), select the Host 2 LSS and volumes in the same way you did it in the previous step. Tivoli Storage Productivity Center for Replication automatically matches all volumes from the selected LSS in Host 1 storage subsystem with all volumes from the selected LSS in Host 2 storage subsystem. In this example, we selected All Volumes. Click Next to continue.



Figure 10-14 Add Copy Set to Metro Mirror session: Host2

- 14. The next window (Figure 10-15) displays a message regarding the matching results. In our example, we have a message that all copy sets matching were successful. However, you might receive warning messages for the following reasons:
 - The number of volumes at Host 1 storage subsystem LSS and the Host 2 storage subsystem LSS is not the same.
 - Volumes at Host 2 storage subsystem LSS are smaller then Host 1 storage subsystem LSS volumes.
 - The Host 1 or Host 2 storage subsystems are already defined in some other copy services session.

The warning messages do not mean the copy sets creation failed. Click **Next** to see the list of available copy sets.



Figure 10-15 Copy set matching results

15. All copy sets volumes that met the matching criteria are automatically selected. You can still modify the current selection and deselect any of the volume pairs included in the list. The Show hyperlink next to each matching volume pair provides copy set information. We selected all copy sets, as shown in Figure 10-16. Click **Next** to continue.

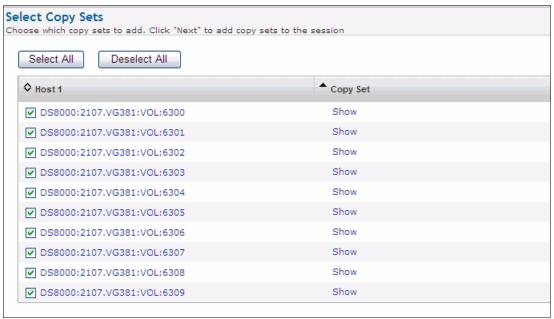


Figure 10-16 Selected copy sets

16. The next window displays the number of copy sets that will be created and the number of unresolved matches (or not selected), as shown in Figure 10-17. Click **Next** to continue.

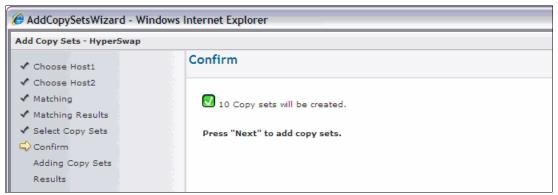


Figure 10-17 Copy set confirmation

17. Tivoli Storage Productivity Center for Replication internally adds that copy set to its database. Figure 10-18 displays a progress panel that reports the number of copy sets added to the Tivoli Storage Productivity Center for Replication inventory database. This step does not start to establish Metro Mirror copy pairs. It is simply a Tivoli Storage Productivity Center for Replication internal process to add this copy set to the Tivoli Storage Productivity Center for Replication database (inventory database).

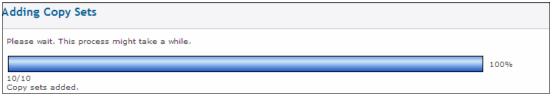


Figure 10-18 Adding Copy Sets

After several seconds, the progress panel reaches 100% and closes. The next panel opens (Figure 10-19). Click **Finish** to exit the Add Copy Sets wizard.



Figure 10-19 Adding Copy Sets results

Figure 10-20 shows and confirms that the session is populated with Copy Sets. The status of the session is still Inactive.



Figure 10-20 HyperSwap session populated with Copy Sets

18. After the session is defined and populated with Copy Sets, you can start the Metro Mirror session with Open HyperSwap. Initially the session can only be started in the direction from Host 1 to Host 2. To start it, select **Start H1->H2** from the Select Action pull-down menu, and then click **Go** (Figure 10-21).

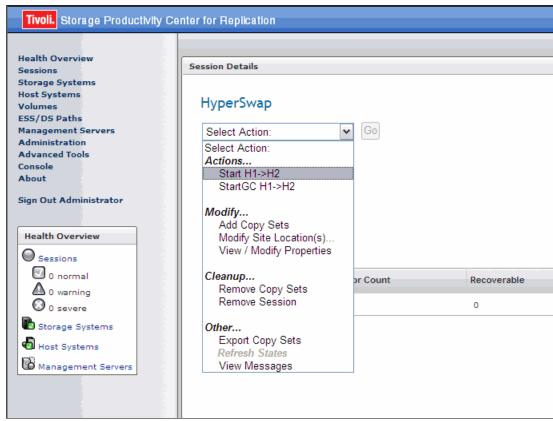


Figure 10-21 Starting the session

19. The next message shown in Figure 10-22 warns that you are about to initiate a Metro Mirror session. It starts copying data from Host 1 to Host 2 volumes, defined previously by adding copy sets, and therefore overwrites any data on Host 2 volumes. Click **Yes** to continue.



Figure 10-22 Session start warning message

In Figure 10-23, the message at the top confirms that the start of Metro Mirror session with Open HyperSwap is completed. The session is in Preparing state and with Warning status.



Figure 10-23 Session start completed

Figure 10-24 shows the Metro Mirror session with Open HyperSwap session without errors. The copy progress is 100% and the session has changed to Prepared state with Normal status.

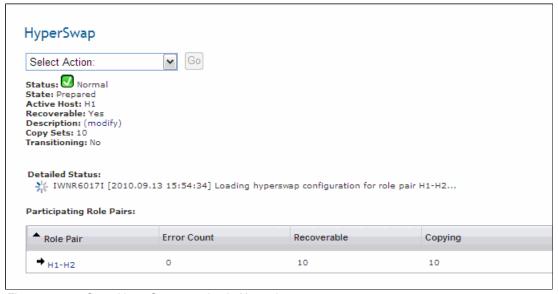


Figure 10-24 Open HyperSwap session in Normal state

Checking the Open HyperSwap session on the AIX host

After the Open HyperSwap configuration has been loaded to the AIX host, the session is ready to perform Open HyperSwap. Example 10-3 shows you the status of AIX hdisks, which are in Open HyperSwap session. The example shows only two AIX hdisks with primary and secondary volumes. The OS direction is from H1 to H2.

Example 10-3 SDDPCM path query device command of Open HyperSwap devices

```
jerome> pcmpath query device
Total Dual Active and Active/Asymmetrc Devices: 10
DEV#:
      5 DEVICE NAME: hdisk5 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME:
              HyperSwap
OS DIRECTION:
              H1->H2
PRIMARY
        SERIAL: 75VG3816300
SECONDARY SERIAL: 75VG4116300
_____
                               State Mode Select
                                                        Errors
Path#
        Adapter/Path Name
                                                0
           fscsi0/path0
                               CLOSE
                                      NORMAL
                                                             0
   0
   1
             fscsi0/path1
                               CLOSE
                                      NORMAL
                                                             0
DEV#: 6 DEVICE NAME: hdisk6 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME: HyperSwap
OS DIRECTION:
              H1->H2
PRIMARY
        SERIAL: 75VG3816301
SECONDARY SERIAL: 75VG4116301
_____
Path#
        Adapter/Path Name
                               State
                                      Mode
                                               Select
                                                         Errors
         fscsi0/path0
                              CLOSE
                                               0
                                                             0
   0
                                      NORMAL
   1
            fscsi0/path1
                               CLOSE
                                      NORMAL
                                                   0
                                                             0
```

You can check the status of Open HyperSwap session using SDDPCM command (Example 10-4).

Example 10-4 SDDPCM query session command

```
jerome> pcmpath query session
Total Open Hyperswap Sessions: 1
SESSION NAME: HyperSwap
SessId
           Host_OS_State
                          Host_copysets
                                         Disabled
                                                       Quies
                                                                Resum
                                                                         SwRes
                   READY
    0
                                                 0
                                                          0
                                                                    0
                                                                             0
```

In our example, we created a new Volume Group and a file system using the Open HyperSwap LUNs. After mounting the file system, we generate I/O activity on it.

As shown in Example 10-5, an asterisk next to the primary volume serial number shows that I/O activity exists on the primary volumes.

Example 10-5 SDDPCM path query device command showing I/O activity on primary volumes

jerome> pcmpath query device Total Dual Active and Active/Asymmetrc Devices: 10 DEV#: 5 DEVICE NAME: hdisk5 TYPE: 2107900 ALGORITHM: Load Balance SESSION NAME: HyperSwap OS DIRECTION: H1->H2 PRIMARY SERIAL: 75VG3816300 * _____ Path# Adapter/Path Name State Mode Select Errors 0 fscsiO/path1 OPEN NORMAL 457 0 SECONDARY SERIAL: 75VG4116300 -----Adapter/Path Name State Mode Select Errors fscsi0/path0 OPEN NORMAL 8 0 Path# 1 DEV#: 6 DEVICE NAME: hdisk6 TYPE: 2107900 ALGORITHM: Load Balance SESSION NAME: HyperSwap OS DIRECTION: H1->H2 ______ PRIMARY SERIAL: 75VG3816301 * _____ Path# Adapter/Path Name State Mode Select Errors 0 fscsi0/path1 OPEN NORMAL 232 0 SECONDARY SERIAL: 75VG4116301 _____ Path# Adapter/Path Name State Mode Select Errors 1 fscsiO/pathO OPEN NORMAL 8 0

10.2.4 Performing Open HyperSwap

After a session has been enabled for Open HyperSwap and reaches the Prepared state, Tivoli Storage Productivity Center for Replication loads the configuration of volumes that are capable of being swapped onto an AIX host. When the load is complete, the session is capable of a planned or unplanned swap.

A planned swap can be used when planned maintenance or a migration from the primary storage is required. Unplanned swap is used in the event of a primary storage system I/O error, and the swap occurs automatically. The I/O switch to the secondary storage system occurs without any user interaction and with minimal application impact. In both cases planned or unplanned the result is the same.

HyperSwap and Open HyperSwap:

- ► For HyperSwap on System z, z/OS completes the entire swap and then alerts Tivoli Storage Productivity Center for Replication that a swap has occurred
- ► For Open HyperSwap, Tivoli Storage Productivity Center for Replication and the AIX host work together to perform the swap.

Planned Open HyperSwap

After a planned swap is run for Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session will have a status of Normal and will indicate that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All I/O should have been redirected to the H2 volumes. After a successful swap to site 2, re-enabling copy to site 2 is not possible. Therefore, issuing a **Start H1->H2** command is not possible. The only way to restart the copy is a **Start H2->H1** command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

Unplanned Open HyperSwap

After an unplanned swap occurs for Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session has a status of Normal and indicates that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All I/O should have been redirected to the H2 volumes. After a successful swap to site 2, re-enabling copy to site 2 is not possible. Therefore, issuing a **Start H1->H2** command is not possible. The only way to restart the copy is by using a **Start H2->H1** command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

Performing a planned Open HyperSwap

In the following example, we show you how to perform a planned Open HyperSwap. The planned Open HyperSwap can be used when planned maintenance or a migration from the primary storage is required.

Complete the following steps:

1. From the Sessions panel, select your Metro Mirror with Open HyperSwap session. From the Select Action pull-down list, select **HyperSwap**, and click **Go** (Figure 10-25).

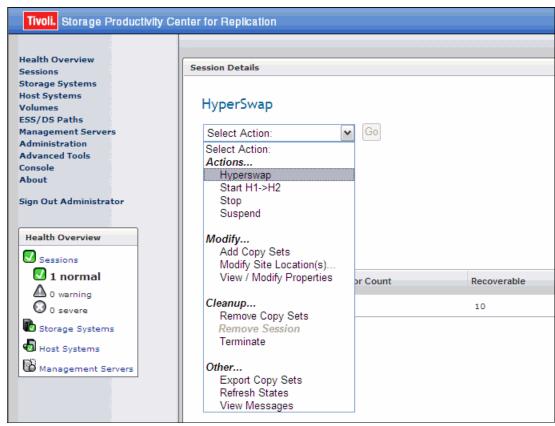


Figure 10-25 Select Open HyperSwap action

2. The next message, shown in Figure 10-26, warns that you are about to initiate an Open HyperSwap action and will move the application I/O from H1 to H2 volumes. Click **Yes** to continue.

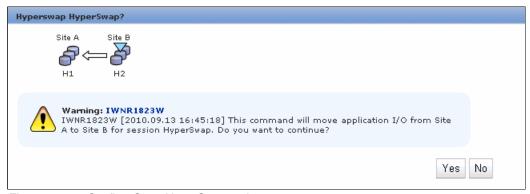


Figure 10-26 Confirm Open HyperSwap action

In Figure 10-27, the message at the top confirms that the Open HyperSwap action is successfully completed. All application I/O are moved to H2 volumes.



Figure 10-27 Open HyperSwap successfully completed

Figure 10-28 shows console log with all actions performed by Tivoli Storage Productivity Center for Replication during the swap.

```
2010.09.13 16:48:07: Administrator: IWNR1028I: The command Hyperswap in session HyperSwap has been run.
2010.09.13 16:48:07: Administrator: IWNR6022I: Issuing Hyperswap to role pair H1-H2...
2010.09.13 16:48:07: Administrator: IWNR6002I: Suspending all pairs in role pair H1-H2...
2010.09.13 16:48:08: Administrator: IWNR1041I: The command freeze was successfully issued to all pairs under role pair H1-H2 for session HyperSwap.
2010.09.13 16:48:08: Administrator: IWNR6020I: Releasing I/O for all pairs in role pair H1-H2...
2010.09.13 16:48:08: Administrator: IWNR1041I: The command thaw was successfully issued to all pairs under role pair H1-H2 for session HyperSwap.
2010.09.13 16:48:08: Server: IWNR1950I: Session HyperSwap changed from the Prepared state to the Suspended state.
2010.09.13 16:48:08: Administrator: IWNR6004I: Recovering all pairs in role pair H1-H2...
2010.09.13 16:48:08: Administrator: IWNR1041I: The command failover was successfully issued to all pairs under role pair H1-H2 for session HyperSwap.
2010.09.13 16:48:08: Server: IWNR1950I: Session HyperSwap changed from the Suspended state to the Target Available state.
2010.09.13 16:48:09: Administrator: IWNR6001I: Stopping all pairs in role pair H1-H2...
2010.09.13 16:48:09: Administrator: IWNR6004I: Recovering all pairs in role pair H1-H2...
2010.09.13 16:48:09: Administrator: IWNR6004I: Recovering all pairs in role pair H1-H2...
2010.09.13 16:48:09: Administrator: IWNR6004I: Recovering all pairs in role pair H1-H2...
2010.09.13 16:48:09: Administrator: IWNR6004I: Recovering all pairs in role pair H1-H2...
2010.09.13 16:48:09: Administrator: IWNR6004I: Recovering all pairs in role pair H1-H2...
2010.09.13 16:48:09: Administrator: IWNR6004I: Recovering all pairs in role pair H1-H2...
```

Figure 10-28 Open HyperSwap console log

3. After a successful swap, re-enabling copy to H2 volumes is not possible. Therefore, issuing a Start H1->H2 command is not possible. The only way to restart the copy is by using a Start H2->H1 command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed, and then the session must be manually restarted to begin copying to the other site. To restart the session from the Sessions panel, select your Metro Mirror with Open HyperSwap session.

From the Select Action pull-down list select **Start H2->H1** and click **Go**, as shown in Figure 10-29.

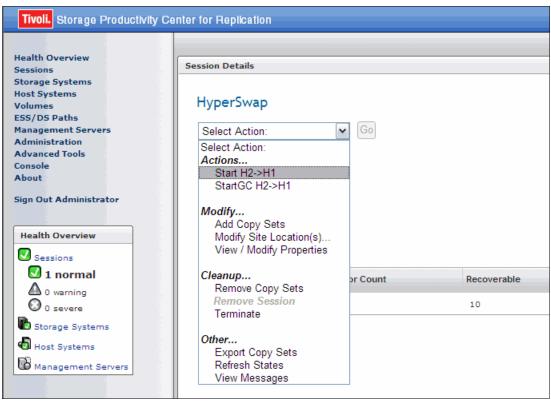


Figure 10-29 Start H2->H1 copy

4. The next message shown in Figure 10-30 is a warning that you are about to initiate copying data from H2 to H1. Click **Yes** to continue.



Figure 10-30 Confirm Start H2->H1 action

In Figure 10-31, the message at the top confirms that the **Start H2->H1** action has successfully completed and that data is copying from H2 to H1.

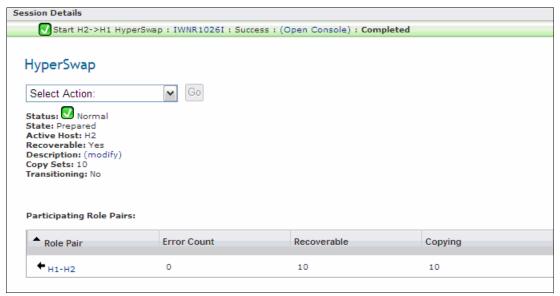


Figure 10-31 Start H2->H1 completed successfully

5. The status of the volumes on the AIX host in Example 10-6 shows that OS direction has changed to H2->H1. It also shows an asterisk (*) next to the secondary volume serial number, which means that the I/O activity is moved to the secondary volumes.

Example 10-6 SDDPCM path query device command showing I/O activity on secondary volumes

jerome> pcmpath query device Total Dual Active and Active/Asymmetrc Devices : 10									
DEV#: 5 DEVICE NAME: hdisk5 TYPE SESSION NAME: HyperSwap OS DIRECTION: H1<-H2	E: 2107900	ALGORITHM	1: Load	Balance					
PRIMARY SERIAL: 75VG3816300									
Path# Adapter/Path Name 0 fscsi0/path1 SECONDARY SERIAL: 75VG4116300 *	State OPEN	Mode NORMAL	Select 457	Errors 0					
Path# Adapter/Path Name 1 fscsi0/path0	State OPEN	Mode NORMAL	Select 8	Errors 0					
DEV#: 6 DEVICE NAME: hdisk6 TYPE: 2107900 ALGORITHM: Load Balance SESSION NAME: HyperSwap OS DIRECTION: H1<-H2									
PRIMARY SERIAL: 75VG3816301									
Path# Adapter/Path Name 0 fscsi0/path1 SECONDARY SERIAL: 75VG4116301 *	State OPEN	Mode NORMAL	Select 232	Errors 0					
Path# Adapter/Path Name 1 fscsi0/path0	State OPEN	Mode NORMAL	Select 8	Errors 0					

Tip: If you want to go back and change OS direction to **H1->H2**, you must swap from H2 to H1 and complete the same steps as described in the example.

10.2.5 Tivoli Storage Productivity Center for Replication high availability with Open HyperSwap

If you have Tivoli Storage Productivity Center for Replication servers in a high availability configuration (active and standby management server configuration) with an Open HyperSwap session, you must be sure that communication is continual between your AIX host and Tivoli Storage Productivity Center for Replication servers.

In an active and standby management server configuration, when you add a host system to the active server the host system connection is automatically sent to the standby server and a connection is attempted. After the communication between the AIX host and the Tivoli Storage Productivity Center for Replication servers is established and the Open HyperSwap session configuration is loaded on the AIX host, your environment is enabled and ready for Open HyperSwap.

In an active and standby management server configuration, the AIX host is using active management server to manage the swap. If a takeover is performed on a standby server, the Open HyperSwap configuration must be loaded from the standby server to the AIX host. After it has been loaded, the AIX host is capable of performing Open HyperSwap actions.

If a takeover is performed on a standby server that is unable to connect to the AIX host system that is managing the swap, the session is no longer capable for Open HyperSwap. Communication to the host system must be activated before the session can become Open HyperSwap capable again.

Attention: When you perform a takeover and use the standby server, ensure that you shut down the active management server first. You must also ensure that you do not have two active management servers. If there are two active management servers and a condition occurs on the storage systems, both management servers respond to the same conditions, which might lead to unexpected behavior.

If you perform an action on the active management server when the servers are disconnected, the servers will be out of synch.

In the following example, we show the takeover action where we have established an AIX host connection to both active and standby management servers.

Figure 10-32 shows the active management server with Open HyperSwap session. In our example, the session is called HyperSwap.

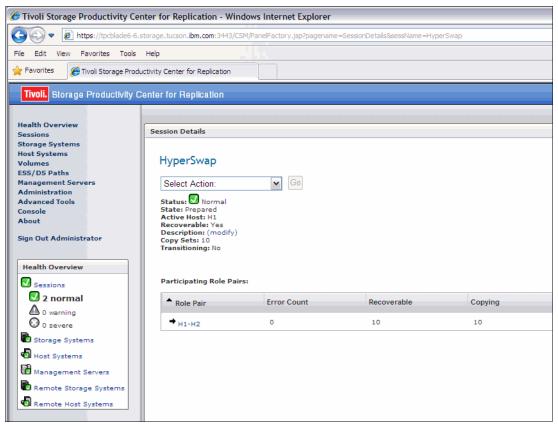


Figure 10-32 Active management server: with HyperSwap session

Takeover action must be started from the standby management server. On the standby management server, select **Management Servers** from the menu on the left.

In the Management Servers panel, select **Takeover** from the drop-down menu (Figure 10-33). Click **Go**.

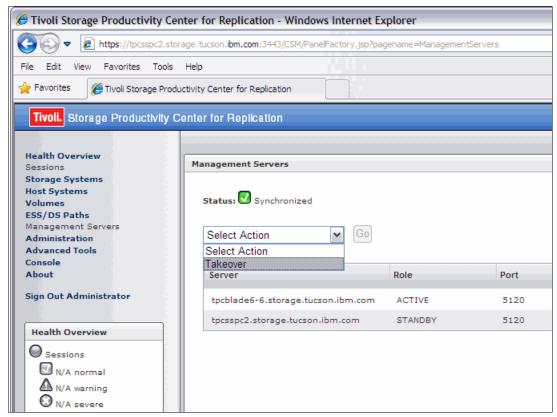


Figure 10-33 Selecting Takeover action

A warning message indicates that both management servers will become active with identical configuration (Figure 10-34). You must ensure that you shut down the active management server first.



Figure 10-34 Takeover warning

After the takeover is successfully completed, the standby management server has the role ACTIVE (Figure 10-35).

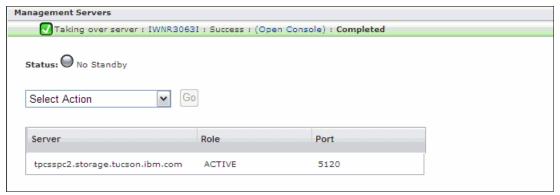


Figure 10-35 Standby server after takeover action

The Open HyperSwap session status changes to severe while the configuration is being loaded to the AIX host (Figure 10-36). In our example, you also see Metro Mirror session without Open HyperSwap. This session status is not changed during the takeover action because it is not dependent on Tivoli Storage Productivity Center for Replication and host connectivity.

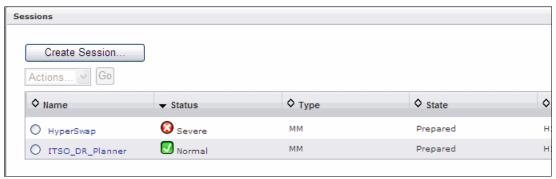


Figure 10-36 Open HyperSwap status: Severe

After the configuration is loaded to the AIX host, the status changes to Normal and the AIX host will be capable of performing Open HyperSwap actions (Figure 10-37).

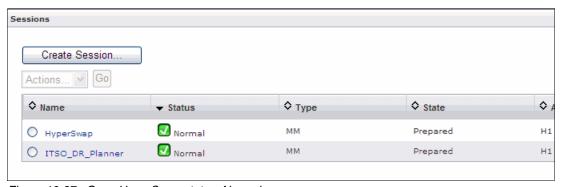


Figure 10-37 Open HyperSwap status: Normal

Figure 10-38 shows the Tivoli Storage Productivity Center for Replication console log of the takeover action.

2010.09.15 10:45:39 : Server : IWNR1951I : High-Availability relationship tpcsspc2.storage.tucson.ibm.com->tpcblade6-6.storage.tucson.ibm.com changed to the Disconnected Consistent state.

2010.09.15 10:45:39 : Administrator : IWNR3086I : High availability has taken over tpcblade6-6.storage.tucson.ibm.com on the server tpcsspc2.storage.tucson.ibm.com.

2010.09.15 10:45:39 : Administrator : IWNR3063I : Successfully issued the takeover to the standby server tpcsspc2.storage.tucson.ibm.com with the active HA server tpcblade6-6.storage.tucson.ibm.com.

2010.09.15 10:45:41 : Server : IWNR6017I : Loading hyperswap configuration for role pair H1-H2...

2010.09.15 10:45:41 : Server : IWNR5501I : Open HyperSwap load completed for session HyperSwap.

Figure 10-38 Console log

10.3 Copy set soft removal of a hardware relationship

Tivoli Storage Productivity Center for Replication V4.2 adds enhancements to the copy set removal function. Before V4.2, this function was used to remove copy set from Tivoli Storage Productivity Center for Replication and also to remove hardware relationships. With V4.2, you can specify, for various reasons, whether to remove a copy set from Tivoli Storage Productivity Center for Replication while leaving the basic hardware relationship.

This enhancement of copy set soft removal can be used for migrations from one session type to another. By using this option you can avoid performing a full copy, which in some scenarios is a time consuming action. Copy set soft removal can be also used for resolving various field problems.

In this section, we show how to remove a copy set while leaving the basic hardware relationship.

Important: Copy set soft removal leaves only base copy relationships for FlashCopy, Global Copy, and Metro Mirror on the storage systems. The relationships will be removed from any consistency groups that are defined on the storage system.

10.3.1 Copy set soft removal scenario

To remove copy set using soft removal of hardware relationship, complete the following steps:

 In the IBM Tivoli Storage Productivity Center for Replication Navigation Tree, select Sessions. Click the radio button next to the session from which you want to remove copy sets, select Remove Copy Sets from the Select Action menu (Figure 10-39), and then click Go. The Remove Copy Set wizard starts.

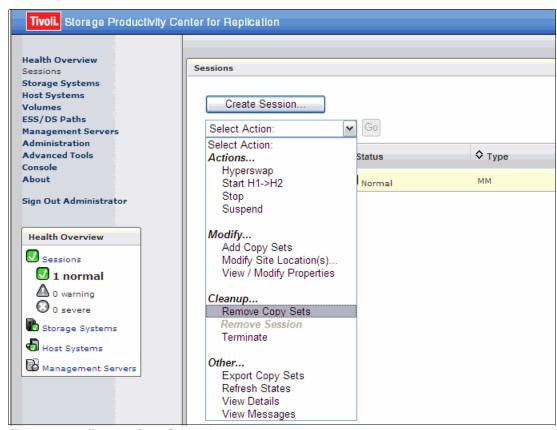


Figure 10-39 Remove Copy Sets action

 From the drop-down menus in the Remove Copy Sets wizard, select the Host 1 storage system, logical storage subsystem, and volume or select the All Volumes option. If you select all for a filter, the lower level filter or filters will be disabled (Figure 10-40). Click Next.



Figure 10-40 Remove Copy Sets - choose Host 1 storage system

3. Select the copy sets that you want to remove and click **Next** (Figure 10-41). In our example, we select only one copy set (volume 6300).

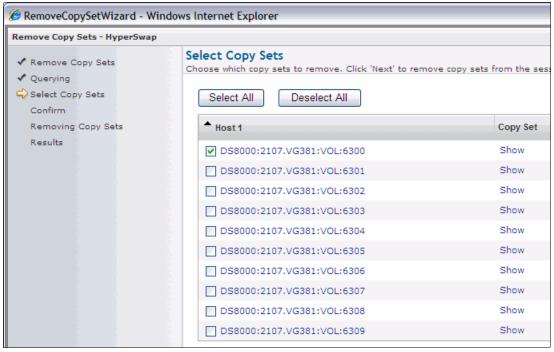


Figure 10-41 Selecting Copy Sets to remove

4. The panel in Figure 10-42 shows the number of selected copy sets to be removed. In this panel, you can select whether to remove or keep the hardware relationship. In our example, we keep the base hardware relationship on the storage system. You can also select to force the copy set to be removed from the session if any errors occurred. The force command forces the removal from the Tivoli Storage Productivity Center for Replication session even if hardware errors occurred. Click **Next** to continue.



Figure 10-42 Selecting the rcopy set removal option

Figure 10-43 shows the progress of removing the copy set.



Figure 10-43 Removing Copy Set process

5. After the process is successfully finished, click **Finish** to exit the wizard (Figure 10-44).



Figure 10-44 Result of removing Copy Sets

6. If you check the DS8000 GUI, you see that the status of the removed copy set is not changed and that it is not removed from the storage system (Figure 10-45). The source and the target volume in our example is 6300.

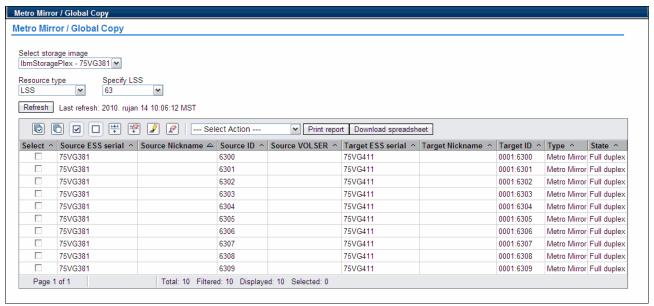


Figure 10-45 Copy set status in DS8000 GUI

10.3.2 Suspending the session after removing a copy set

After you remove a copy set from a session, and if you suspend the session, all copy sets within the session will be suspended, including the copy set that you previously removed. Suspend action is done on the LSS level, so all volumes within the same LSS will also be suspended.

In our example, we removed one copy set (volume 6300) and we left the hardware relationship as shown in Figure 10-45. The session now contains nine copy sets (Figure 10-46).

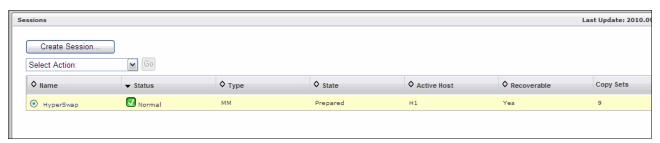


Figure 10-46 Session status

After we suspend the session (Figure 10-47), the removed copy set (volume 6300) will also be suspended. Figure 10-48 shows you the status of copy sets in DS8000 GUI after suspend action.



Figure 10-47 Session suspended

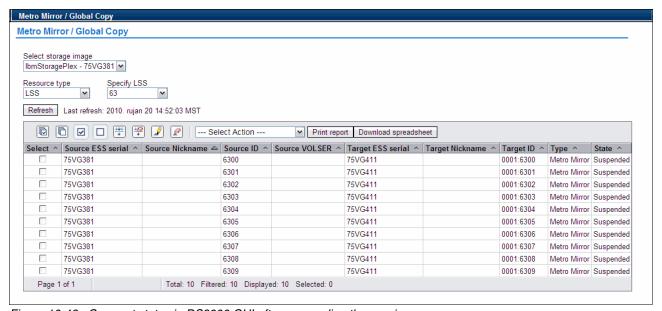


Figure 10-48 Copy set status in DS8000 GUI after suspending the session

If you resume the session, only the copy sets within the session are resumed, while the removed copy set remains in suspended status (Figure 10-49).

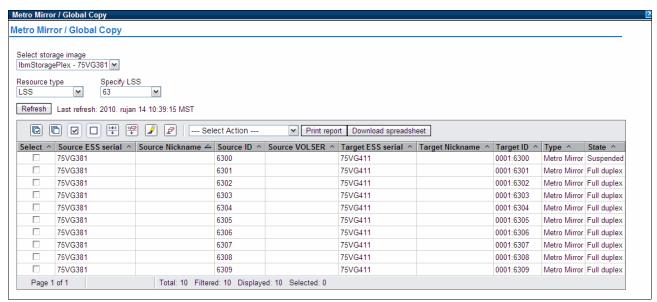


Figure 10-49 Copy set status in DS8000 GUI after resuming the session

10.3.3 Removing a copy set from an Open HyperSwap

If you remove a copy set from an Open HyperSwap session, the copy set can remain coupled on the host. It happens if an application on the host has opened a device and the copy set in a session for that device is removed. Figure 10-50 shows you the error from Tivoli Storage Productivity Center for Replication console.

Tip: Because of this error, always close the application and then remove copy set, so that remove copy set can decouple the devices.



Figure 10-50 Coupled copy set error

Example 10-7 shows the SDDPCM **pcmpath query device** command for a coupled copy set. The session name is blank, because the device is no longer associated with the session.

Example 10-7 SDDPCM pcmpath query device command for a coupled copy set

The general steps to decouple the copy set are as follows:

- 1. On the AIX host, stop the application that has opened the device, unmount file system, and vary offline the volume group.
- 2. Remove the device by using the following command, where hdisk_number is the number of the hdisk (device) that you want to remove:
 - rmdev -dl hdisk hdisk number
- Run the cfgmgr command to discover the device. If you run pcmpath query device command, separate devices are presented for the copy set pair.

10.4 Downloading log packages with the GUI

In Tivoli Storage Productivity Center for Replication V4.2, you can download log packages without the access to the Tivoli Storage Productivity Center for Replication file system. The log packages can help you more quickly with diagnostics of IBM Tivoli Storage Productivity Center for Replication anomalies.

You can download a log package to the local system from the GUI as follows:

 Log on to Tivoli Storage Productivity Center for Replication. From the Tivoli Storage Productivity Center for Replication menu, select **Advanced Tools**, which opens the panel in Figure 10-51.

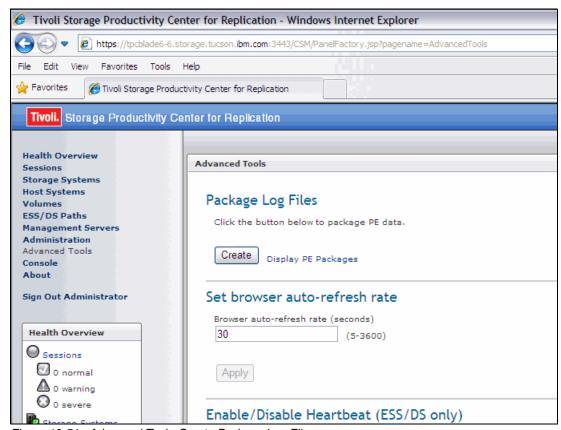


Figure 10-51 Advanced Tool:- Create Package Log Files

2. Click **Create**. The packages start being created (Figure 10-52).

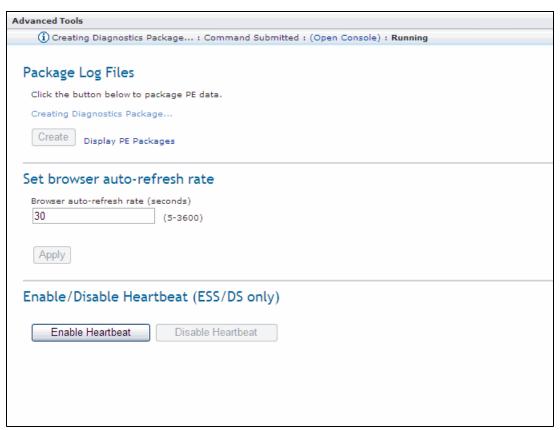


Figure 10-52 Creating Diagnostic Packages

3. After the packages are created, a success message is issued (Figure 10-53).

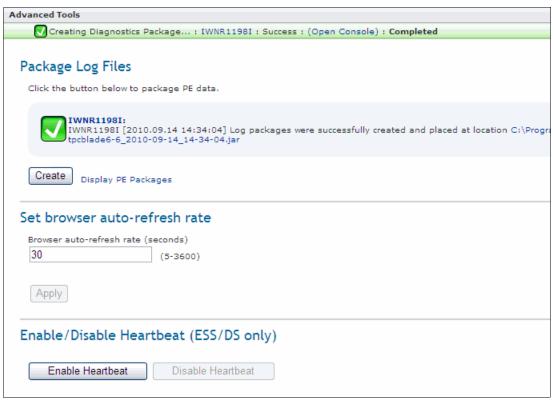


Figure 10-53 Diagnostic Packages created successful

4. Click the hyperlink **Display PE Packages**. The new or existing packages are listed (Figure 10-54). Click the package name and save it locally to your selected folder.



Figure 10-54 Diagnostic Log Packages

10.5 Path Manager

Path Manager is a new function in Tivoli Storage Productivity Center for Replication V4.2 and provides control of logical paths when establishing relationships between source and target storage systems.

The logical paths define the relationship between a source logical subsystem (LSS) and a target LSS that is created over a physical path. Path Manager can help you to explicitly control which port pairing Tivoli Storage Productivity Center for Replication will use when establishing the logical paths, and ensure redundant port combinations. It also keeps that information persistent for use when the path is terminated because of a suspend operation.

Tivoli Storage Productivity Center for Replication V4.2 provides you with the following options to create the logical paths and specify port pairing:

- ► Adding logical paths automatically. Tivoli Storage Productivity Center for Replication automatically picks the paths or uses established paths.
- Adding logical paths and creating port pairing by using a comma separated (CSV) file
- Adding logical paths by using Tivoli Storage Productivity Center for Replication GUI

Attention: The Path Manager does not affect, modify, or deal with Global Mirror control paths. It applies only to data paths between Metro Mirror or Global Copy relationships.

Adding logical paths automatically

If you do not specify the port pairing, Tivoli Storage Productivity Center for Replication will query the storage systems and automatically pick the available ports to create the paths. If you have established paths, Tivoli Storage Productivity Center for Replication will use those paths and it will not create the new one.

This option does not guarantee that you will have redundant logical paths.

Adding logical paths using a CSV file

You can add logical paths by creating a comma-separated values (CSV) file. The CSV file specifies storage systems pairings and associated port pairings that will be used by Tivoli Storage Productivity Center for Replication to establish the logical paths. By using the CSV file, you can ensure redundant port combinations and use only the specified ports. Tivoli Storage Productivity Center for Replication uses the ports listed in the CSV file when you issue a start command (that is, **Start H1->H2**) and it attempts to establish the paths between any LSS on those two storage systems.

Using the CSV file:

- ► Use the CSV file to establish the logical paths if you have a large environment (large number of LSS). By using the CSV file, managing port pairings can be easier because you do not have to specify ports for each LSS in relationship.
- ▶ If you use the CSV file, Tivoli Storage Productivity Center for Replication always uses the specified ports that ensure redundant port combination.
- ▶ If you use the CSV file, always manage your port pairings by using only the CSV file.

Steps for adding logical paths using a CSV file

To add logical paths using a CSV file, complete the following steps:

1. Create a CSV file named portpairings.csv in the following directory:

install_root/eWAS/profiles/ CSM/properties

Tip: The install_root/eWAS/profiles/ CSM/properties directory contains the portpairings.csv.sample file, which you can use and rename to portpairings.csv file.

You can also create the CSV file in a spreadsheet such as Microsoft Excel or in a text editor. An example of a CSV file is shown in Example 10-8.

Example 10-8 Port pairings CSV sample file

```
# Examples:
2107.04131:2107.01532,0x0331:0x0024,0x0330:0x0100,0x0331:0x000C
2107.05131:2107.01532,0x0330:0x0029,0x0331:0x0001
```

Each line in the file represents a storage system to storage system pairing. The first value represents the storage systems, which are delimited by a colon. The remaining values are the port pairs, which are delimited by a colon. All values are separated by a comma and commented lines must start with #.

Tip: When specifying a storage system in the CSV file, use only the last five numbers from the storage ID. For example, if the storage ID is IBM.2107-7516381, specify 2107.16381 as the storage system.

To enable the changes in the file, use a task that requires new paths to be established. For example, suspend a session to remove the logical paths and then issue the Start H1->H2 command to enable the paths to use the port pairings in the CSV file. After the start command finishes successfully, the paths are listed as No in the Auto Generated column in the ESS/DS panel (Figure 10-55).

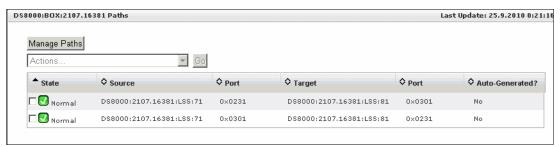


Figure 10-55 Successfully established paths using CSV file

Rules for using a CSV file

When you use the CSV file, consider the following important rules:

- ► The entry for a storage system pair and the port pairs are bidirectional. This means that a line that has systemA:systemB is equivalent to a line that has systemB:systemA.
- ► Lines that are incorrectly formatted are discarded. For example, if a line contains ports without the 0x, or does not contain port pairs delimited by the colon (:) character, the whole line is discarded.
- ► A line can be properly formatted but contain invalid ports for your given storage system configuration. In this case, the ports will be passed down to the storage system to be established, and no validation is done in Tivoli Storage Productivity Center for Replication. The valid ports might be established by the storage system, and the invalid ones could be rejected. If no valid ports exist, an error occurs (Figure 10-56).
- ► If a file contains duplicate lines for the same storage systems, the ports on the last line are the ones that are used. Entries are bidirectional also. Therefore, if you have systemA:systemB and a line with systemB:systemA, the second line is the line that is used.
- ► Any line that starts with a number sign (#) character is counted as a comment and is discarded. The # character must be at the start of the line. Placing it in other positions can cause the line to be invalid.
- ► The portpairings.csv file is not shared between two Tivoli Storage Productivity Center for Replication servers in a high availability environment. Thus, it is possible that different port pairings can be established from the standby server after a takeover. You have to copy the portpairings.csv file to the standby server to ensure that the two files are equal.

```
24.9.2010 23:58:06: tpc42inst: IWNR1027E: The command Start H1->H2 in session MM has completed with one or more errors.
24.9.2010 23:58:05: tpc42inst: IWNR2108E: A hardware error occurred during the running of a command for the pair in session MM for copy set

50 DS8000:2107.16381:VOL:7100 with source DS8000:2107.16381:VOL:7100 and target DS8000:2107.16381:VOL:8100 in role pair H1-H2. The hardware returned an error code of <a href="https://docs.org/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/linearing/lin
```

Figure 10-56 Error when adding an invalid port

Adding logical paths by using Tivoli Storage Productivity Center for Replication GUI

You can add logical paths using the Tivoli Storage Productivity Center for Replication GUI. This section shows how to add the logical paths. Before you add the paths, ensure that you have defined the appropriate storage systems on the Storage Systems panel.

To add the logical paths, complete the following steps:

1. In the Tivoli Storage Productivity Center for Replication Navigation Tree, select **ESS/DS Paths**. The next panel for managing paths opens (Figure 10-57). Click **Manage Paths**.



Figure 10-57 Manage paths

2. The Path Management wizard opens (Figure 10-58). From the drop-down boxes in the Path Management wizard, select the source storage system, source logical storage system, target storage system, target logical storage system, and then click **Next**.



Figure 10-58 Selecting source and target storage systems

3. From next the drop-down boxes in the Path Management wizard, select the source port and target port, and then click **Add** (Figure 10-59). You can add multiple paths between the logical storage subsystems, or only one at a time.

After making your selections, click Next.

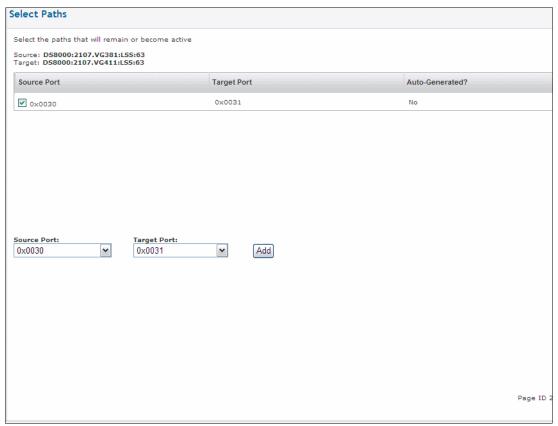


Figure 10-59 Adding ports

4. Confirm your selections and click **Next** (Figure 10-60).

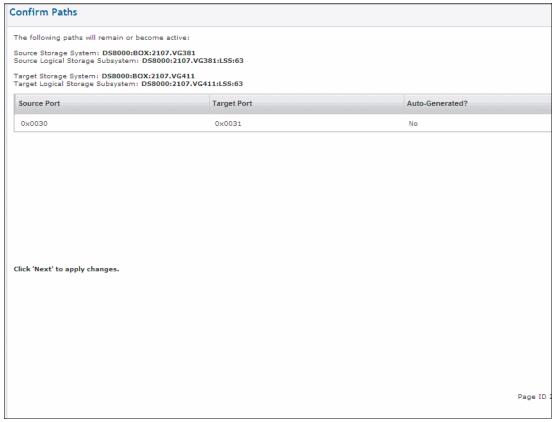


Figure 10-60 Confirm the paths

5. Verify the Results panel and click **Finish** to exit the wizard (Figure 10-61).



Figure 10-61 Results panel

6. Click the Storage System to see the path that you have just added (Figure 10-62).

Figure 10-62 Managed paths

10.6 SVC and Storwize V7000 enhancements

In this section, we describe the SAN Volume Controller (SVC) and Storwize V7000 enhancements in Tivoli Storage Productivity Center for Replication V4.2.

Support: Support for Storwize V7000 is available starting with Tivoli Storage Productivity Center V4.2.1.

10.6.1 SVC and Storwize V7000 space-efficient volumes

Tivoli Storage Productivity Center for Replication V4.2 added support for SVC and Storwize V7000 space-efficient volumes in all Tivoli Storage Productivity Center for Replication SVC and Storwize V7000 sessions.

SVC and Storwize V7000 space-efficient volumes can be used in any role in a Tivoli Storage Productivity Center for Replication session. However, because the space-efficient volumes require additional I/O operations to read and write metadata to a back-end storage, and because they generate additional load on the SVC or Storwize V7000 nodes, be sure to analyze the environment where the space-efficient volumes will be used. It is best not to use space-efficient volumes for high-performance applications.

If you use the space-efficient volumes that are not in Target or Journal roles, a warning is issued (Figure 10-63).

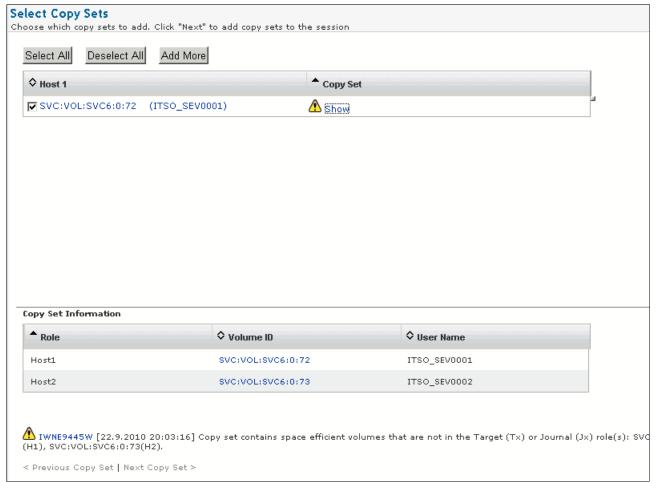


Figure 10-63 Space-efficient volumes warning

Figure 10-64 shows the volumes panel where you can see your SAN Volume Controller space-efficient volumes, which are indicated by "Yes" in the Space Efficient column.

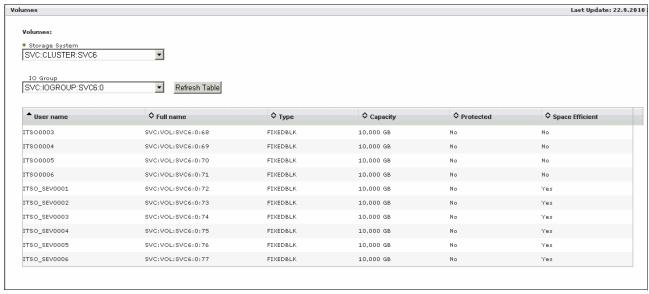


Figure 10-64 SVC space-efficient volumes in Tivoli Storage Productivity Center for Replication

10.6.2 SVC and Storwize V7000 incremental FlashCopy

Tivoli Storage Productivity Center for Replication V4.2 supports SAN Volume Controller and Storwize V7000 incremental FlashCopy in the Tivoli Storage Productivity Center for Replication FlashCopy, Metro Mirror with Practice, and Global Mirror with Practice sessions (Figure 10-65).

FlashCopy: By default, the incremental FlashCopy is enabled in practice sessions and disabled in FlashCopy sessions.

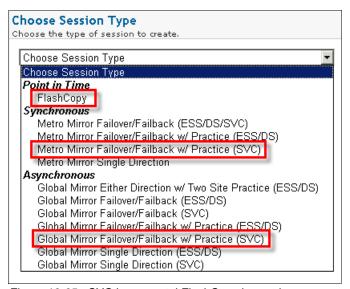


Figure 10-65 SVC incremental FlashCopy in sessions

If you select the incremental option on those sessions, the amount of time to perform a copy operation is significantly reduced. A session that is created with the incremental FlashCopy option will copy only the data that has been changed on the source or the target since the previous copy completed. The incremental FlashCopy can substantially reduce the time that is required to re-create an independent image.

Another example where incremental FlashCopy option can benefit is when the FlashCopy mapping was stopped before the background copy completes and when the mapping is restarted, the data that had been copied before the mapping was stopped is not copied again. For instance, if an incremental mapping reaches 10% progress when it is stopped and restarted, that 10% of data is not re-copied when the mapping is restarted, assuming of course that it was not changed.

Tip: Even if you use the incremental FlashCopy option, the first copy process copies all of the data from the source to the target SVC or Storwize V7000 VDisk.

The following example shows how to create and start an incremental FlashCopy session:

1. In the Tivoli Storage Productivity Center for Replication Navigation Tree, select **Sessions**. Click **Create Session** to open the Create Session wizard (Figure 10-66).



Figure 10-66 Create session wizard

From the drop-down menu, select the FlashCopy session type and click Next (Figure 10-67).



Figure 10-67 Select FlashCopy session

- 3. In the Properties panel (Figure 10-68), type a session name and description. Here you can specify whether you want to use incremental FlashCopy. By selecting this option, you set up the relationship for recording changes to the source volume (H1 volume). This means that any subsequent FlashCopy operation for that session copies only the tracks that have changed since the last flash. Incremental always assumes persistence.
 - In the Background Copy Rate box, you can type the copy rate that the SAN Volume Controller or Storwize V7000 will use to perform the background copy of the FlashCopy role pair. You can specify a percentage in the range of 0 100. The default is 50.

Specifying 0 is equivalent to specifying the No Copy option for a System Storage DS8000 or TotalStorage Enterprise Storage Server FlashCopy session. You can modify this value at any time during the session. If the session is performing a background copy when you change the option, Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency group on the SAN Volume Controller or Storwize V7000. The SAN Volume Controller or Storwize V7000 consistency group begins to use this new rate to complete the background copy that it is performing. Click **Next** to continue.

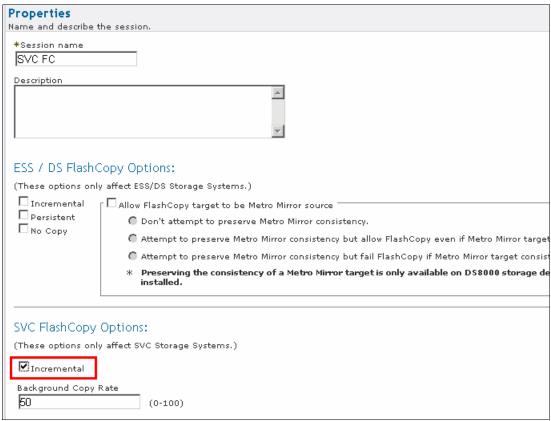


Figure 10-68 Selecting FlashCopy incremental option

4. From the drop-down menu, choose a location for Site 1 and click **Next** (Figure 10-69).

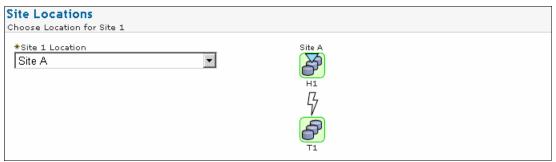


Figure 10-69 Site locations

5. In the next panel, verify that the session was added successfully. If the session is successfully created, click **Launch Add Copy Sets Wizard** to add copy sets to the session. The wizard opens (Figure 10-70).

From the drop-down menus, select the Host 1 storage system, IO group, and volume. For the IO group and volume, you can select all entries, or an individual entry. In our example, we select only one volume from an IO group. If you want to import a copy set, select the Use a CSV file to import copy sets check box. You can either enter the full path name of the CSV file in the text box, or click **Browse** to find the CSV file. Click **Next** to continue.

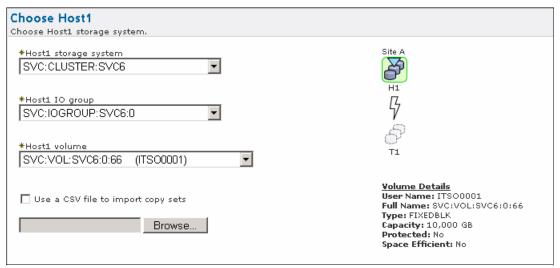


Figure 10-70 Selecting a host storage system

6. From the drop-down menus (Figure 10-71), select the Target 1 storage system, IO group, and volume. In our example, we use space-efficient target volume. Click **Next** to continue.

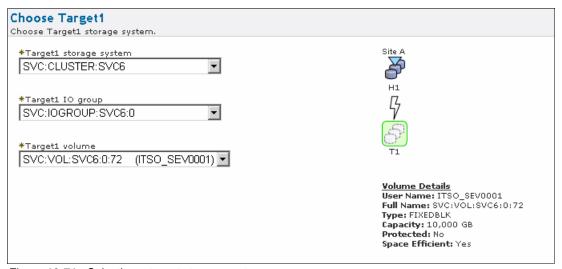


Figure 10-71 Selecting a target storage system

7. The following panel (Figure 10-72) shows you copy sets, which you can select. Select the copy sets that you want to add. You can click **Select All** to select all the boxes, **Deselect All** to clear all the boxes, or **Add More** to add more copy sets to this session. In our example, we add only one copy set.

If you click the **Show** link, the copy set information volumes that you have selected are listed. Click **Next** to continue.



Figure 10-72 Selecting the copy sets

8. Figure 10-73 shows the panel with the number of copy sets that will be added. Click **Next** to confirm and add copy sets. A progress bar displays the status of adding the copy set.



Figure 10-73 Confirm the copy sets

9. When the progress completes, the Results panel opens (Figure 10-74). Click **Finish** to exit the wizard.



Figure 10-74 Adding copy sets result

10. Figure 10-75 shows a defined SVC FlashCopy session with one copy set. The session is in a Defined status and it can be started.



Figure 10-75 SVC FlashCopy session

11.To start the FlashCopy session, click Select Action → Start, and then click Go (Figure 10-76).

Attention: Start action performs any steps necessary to define the relationship before doing a FlashCopy operation. The session has Prepared status.

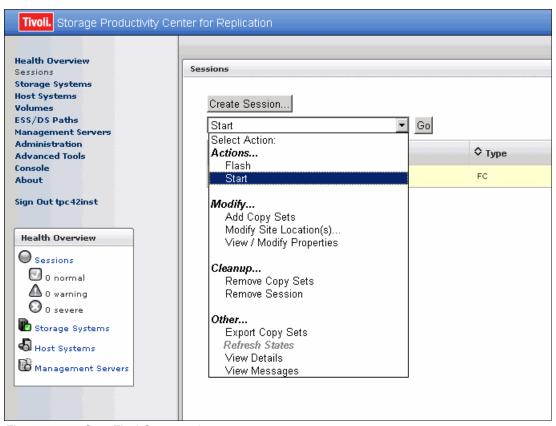


Figure 10-76 Start FlashCopy session

12.A warning message indicates that the relationship will be established and prepared. Click **Yes** to establish the incremental FlashCopy session (Figure 10-77).

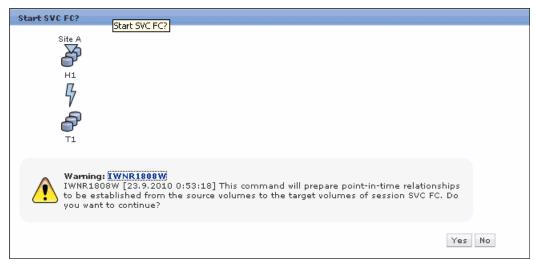


Figure 10-77 Warning before establishing SVC FlashCopy session

After the command completes successfully, the session will be in preparing status and it will change to Prepared status (Figure 10-78 on page 400).

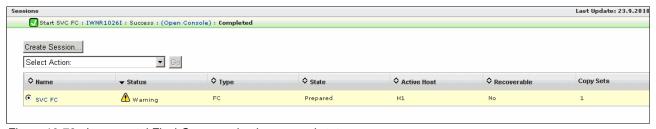


Figure 10-78 Incremental FlashCopy session in prepared status

13.To start the FlashCopy operation, select the session (in the session panel), select the Flash action (from the drop-down menu), and then click Go (Figure 10-79). The Flash action creates a data consistent point-in-time copy. The FlashCopy session with the incremental option ensures that only the regions of disk space where data has been changed since the FlashCopy mapping was last started are copied to the target volume.

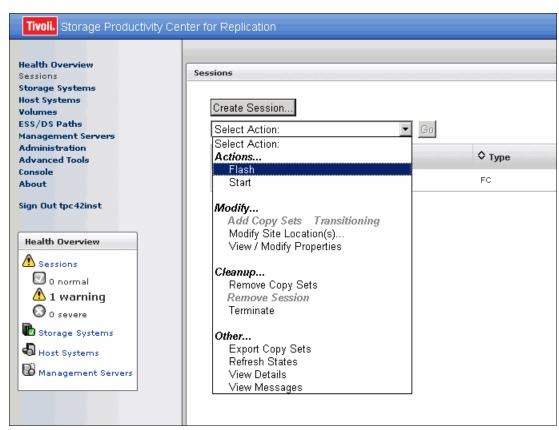


Figure 10-79 Perform FlashCopy operation

14.A warning message indicates that the point-in-time copy will be created (Figure 10-80 on page 401). Click **Yes** to create the incremental FlashCopy.

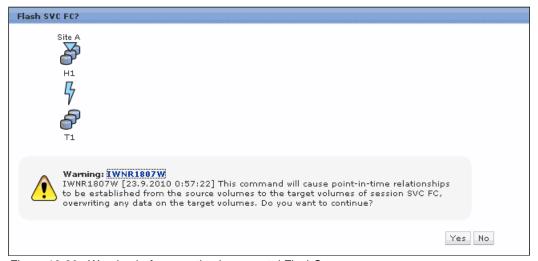


Figure 10-80 Warning before creating incremental FlashCopy

After the command completes successfully (Figure 10-81), the session will be in a Target Available status, and the copy process will start in the background (Figure 10-82).

If, for example, the FlashCopy session is stopped before the background copy completes, when the session gets restarted the data which had been copied before the session was stopped does not get copied again.



Figure 10-81 Successfully created incremental FlashCopy



Figure 10-82 Incremental FlashCopy background process

10.7 DS8000 enhancements

In this section, we discuss the Tivoli Storage Productivity Center for Replication V4.2 features related to DS8000.

10.7.1 DS8000 extent space efficient volumes

Beginning with the DS8000 release 4.3, DS8000 storage systems support two types of space efficient volumes:

- Track space efficient (TSE) volumes
- ► Extent space efficient (ESE) volumes

TSE logical volumes are used as target volumes of a FlashCopy SE operation (with a nocopy option). Because in a nocopy relation, only modified data is copied to the target volumes, this kind of space efficient volume can operate with a high over-provisioning ratio.

ESE volumes are implemented to allow thin provisioning in general and are dynamically provisioned using the DS8000 regular extents. When an ESE logical volume is initially created, the volume has no real data capacity. However, the DS8000 will use some real capacity for metadata that it uses to manage space allocation. The metadata holds information about extents and volumes blocks already allocated in a rank. This metadata is used for thin provisioning, so that the DS8000 can determine whether new extents have to be allocated.

Reference: For the support of space efficient volumes on DS8000 storage systems and required microcode level, see the following addresses for more information:

► DS8100 and DS8300:

http://www.ibm.com/support/docview.wss?rs=1113&uid=ssg1S1002949

► DS8700:

http://www.ibm.com/support/docview.wss?uid=ssg1S1003593&rs=1329

Tivoli Storage Productivity Center for Replication V4.2 added support for ESE volumes, and indicates whether a volume is ESE. It differentiates between ESE and TSE volumes. Figure 10-83 shows the panel with ESE volumes. Figure 10-84 and Figure 10-85 show TSE volumes and standard DS8000 volumes.

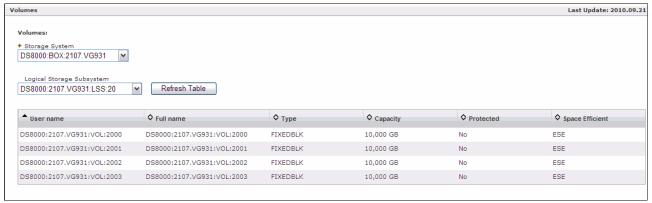


Figure 10-83 Tivoli Storage Productivity Center for Replication showing ESE volumes

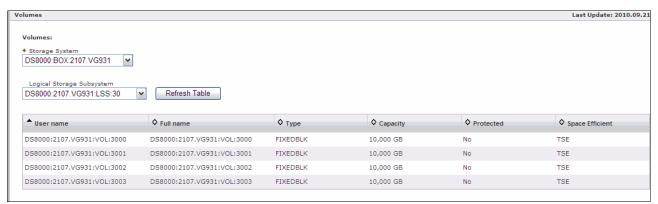


Figure 10-84 Tivoli Storage Productivity Center for Replication showing TSE volumes

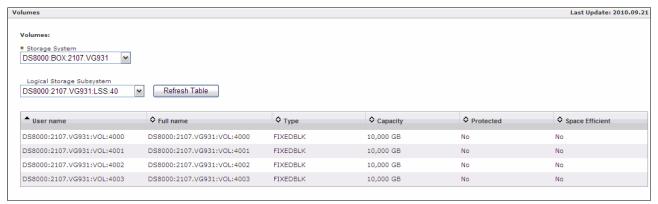


Figure 10-85 Tivoli Storage Productivity Center for Replication showing standard volumes

Even the Tivoli Storage Productivity Center for Replication shows extent space efficient volumes, the ESE volumes can not be used in copy service relationships. If you try to add ESE volumes to a copy set you will get an error. Figure 10-86 shows you the error.

Tip: ESE volumes are currently restricted from participating in copy service relationships.



Figure 10-86 Adding ESE volumes to a copy set

TSE can be used in a copy sets but only in Target or Journal roles. If you try to add it into any other role, you will not be able to select it and the No valid selection message is issued (Figure 10-87).



Figure 10-87 Adding TSE volumes to a copy set

10.7.2 Global Mirror session enhancements

Tivoli Storage Productivity Center for Replication V4.2 provides Global Mirror session enhancements by adding more Global Mirror session information. An important consideration about the Global Mirror session is data exposure, which is provided in a convenient graph, which ensures that data exposure is within a threshold for at least 24 hours. Global Mirror session information enhancements are as follows:

- ▶ Data exposure information
- Requested consistency group interval time
- ► Number of successful consistency groups in a last interval
- Unsuccessful consistency groups during the last formation
- Consistency group interval time
- ► Maximum coordination time
- Maximum consistency group drain time

In the following figures, we show several examples of the new information about the Global Mirror session.

In the session details panel, if you click the Global Mirror role pair Global Mirror session information is displayed. Figure 10-88 shows the new information about the Global Mirror session.

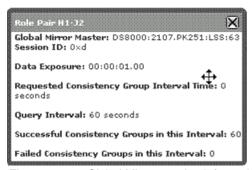


Figure 10-88 Global Mirror session information view

If you click the Global Mirror information panel tab on session details page, the following information is listed (Figure 10-89):

- ► Global Mirror master LSS
- Master consistency group time
- Master time during last query
- ► Data exposure time
- Session ID
- ► Master State
- Unsuccessful Copy Groups during last formation
- ► CG interval time
- ► Maximum Coordination time
- Maximum Copy Group drain time
- List of subordinates (displayed only if they exist)

The table in the panel displays the ratio of successful and unsuccessful Copy Groups formed since the last query and overall.

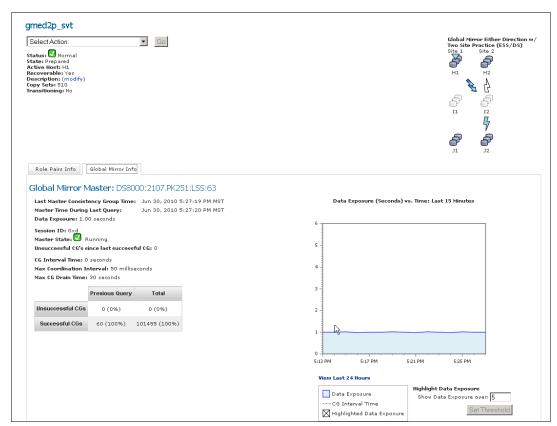


Figure 10-89 GM session information

If unsuccessful Copy Group formations occurred, reason codes are given in expandable sections (Figure 10-90).

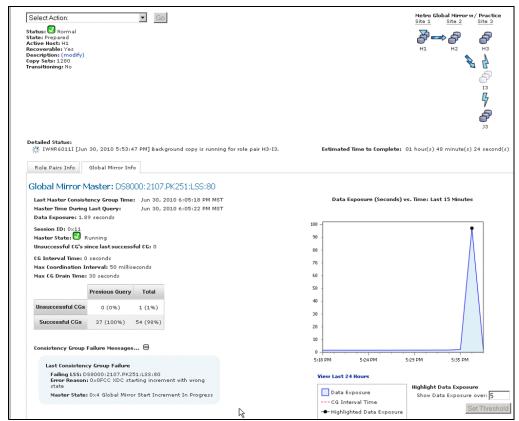


Figure 10-90 Unsuccessful CG formation

The failure message also contains the following information:

- ▶ Failing LSS
- ► Error reason
- Master state

Global Mirror Data Exposure Graph in the next figures shows you the last 15 minutes and last 24 hours interval. It displays the data exposure over time. You can set up a data exposure threshold, which highlight outliers or unusual spikes. The graph also shows consistency group interval time, and if no data is collected because of loss of communication with the storage system, it is also indicated. Figure 10-91 shows data exposure graph for the last 24 hours.

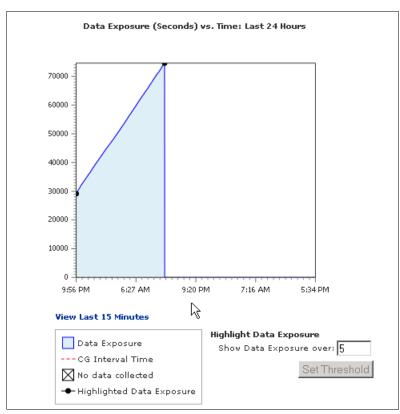


Figure 10-91 GM data exposure during last 24 hours

Figure 10-92 shows the graph for the last 15 minutes.

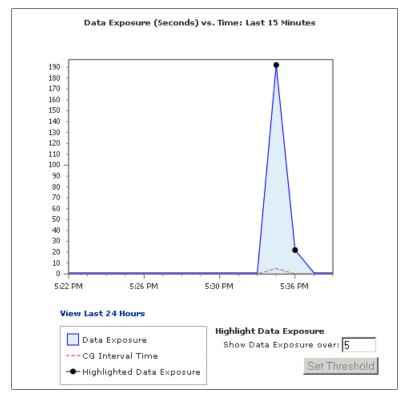


Figure 10-92 GM data exposure during last 15 minutes

10.7.3 Multiple Global Mirror sessions

Tivoli Storage Productivity Center for Replication 4.2 supports multiple Global Mirror sessions functions, allowing you to create more than one Global Mirror session per storage system. Previously, you could have only one Global Mirror session per storage system box, a limitation set by hardware.

Support: The multiple Global Mirror sessions function is available on DS8000 storage systems at request for price quotation (RPQ) with microcode level 5.1. On SVC hardware types, it is already supported. It is included for Storwize V7000 and XIV with their respective updates in Tivoli Storage Productivity Center for Replication 4.2.1 and 4.2.2.

With Tivoli Storage Productivity Center for Replication V4.2, this function allows you to create multiple Global Mirror sessions and enables you to fail over only data that is assigned to one host or application instead of forcing you to fail over all data if one host or application fails. This provides increased flexibility to control the scope of a failover operation and to assign separate options and attributes to each session.

Sessions: The multiple Global Mirror session limitations are as follows:

- Only one Global Mirror session is allowed per LSS.
- ► A maximum of 255 Global Mirror master sessions is allowed per box.
- ► A maximum of 255 Global Mirror subordinate sessions is allowed per box.

Installation and planning

If you are planning to create multiple Global Mirror sessions, you can create them in the same way as with the previous Tivoli Storage Productivity Center for Replication releases. The difference with the Multiple Global Mirror session is that when creating a session, it uses the first available Global Mirror session ID. The Global Mirror session number is chosen (in the background) by an algorithm that avoids already used Global Mirror session IDs. If the limit of Global Mirror sessions has been reached a message is issued.

You can create Global Mirror sessions by using the Tivoli Storage Productivity Center for Replication GUI Create Session wizard or by using the **mksess** CLI command. The general steps to create a Global Mirror session and add copy sets are as follows:

- 1. Create Global Mirror session.
- 2. Add copy sets on a DS8000 storage system.
- 3. Start the Global Mirror session.

Migration

If you are upgrading to Tivoli Storage Productivity Center for Replication V4.2 from the previous Tivoli Storage Productivity Center for Replication releases, a Global Mirror session is maintained during the upgrade. You do not need to take additional steps to migrate your existing sessions.

If you are upgrading Tivoli Storage Productivity Center for Replication with your Global Mirror session and if you have Tivoli Storage Productivity Center for Replication in a high availability configuration, the general steps for the upgrade are as follows:

- 1. Issue a takeover from standby Tivoli Storage Productivity Center for Replication server.
- 2. Upgrade the active Tivoli Storage Productivity Center for Replication server.
- 3. Ensure sessions are all maintained on the active Tivoli Storage Productivity Center for Replication server.
- 4. Upgrade standby Tivoli Storage Productivity Center for Replication server.

10.8 XIV enhancements

In this section, we discuss the Tivoli Storage Productivity Center for Replication V4.2 features related to XIV.

Support: Support for XIV is available starting with Tivoli Storage Productivity Center for Replication V4.2.2. This is only for XIV Gen2 hardware. Look for XIV Gen3 hardware support in a future Tivoli Storage Productivity Center fix pack.

10.8.1 Overview

You can now use Tivoli Storage Productivity Center for Replication 4.2 to manage copy services solutions on XIV storage systems. XIV support is included as a part of Tivoli Storage Productivity Center for Replication and does not require any additional installation. Three new session types have been added for XIV storage systems:

► Snapshot:

Available with all Tivoli Storage Productivity Center editions

Metro Mirror Failover/Failback:

Available with Tivoli Storage Productivity Center for Replication Two Site Business Continuity

► Global Mirror Failover/Failback:

Available with Tivoli Storage Productivity Center for Replication Two Site Business Continuity

There are no practice session types available for XIV storage systems.

Interface: Tivoli Storage Productivity Center for Replication V4.2.2 provides the user with a common interface for managing DS8000, SAN Volume Controller, Storwize V7000, and XIV storage systems.

XIV concepts

The following terms are referenced throughout this chapter with respect to XIV storage systems:

Consistency group A set of volumes that is treated as one

Mirror A replica of a volume or consistency group to another volume or

consistency group.

Pool An allocation of space used to create volumes

Snapshot A point-in-time copy of a given volume or consistency group

Snapshot group A group of snapshots that is formed from a consistency group

On XIV storage systems, primary and secondary volumes are referred to as master and slave volumes. For details on XIV Storage System Copy Services, see IBM Redbooks Publication XIV Storage System -Copy Services and Migration (SG24-7759).

XIV connections

To add a connection to an XIV storage system, select **Storage Systems** → **Add Storage Connection**. This launches a wizard for adding a storage connection for all of the supported storage system device types. When you select **XIV**, you will be prompted to enter connection details, which include a single IP address to the XIV storage system, a user name, and a password (Figure 10-93). The XIV user must have storage administrator authority.



Figure 10-93 Add a single IP address for the XIV connection

Tip: When you add the IP address of an XIV storage system, you get three connections between it and Tivoli Storage Productivity Center for Replication. You do not need to enter all three IP addresses yourself.

After you have completed the wizard to add the storage system, an SSH connection is established, and the XIV shows up in the list of connections (Figure 10-94).

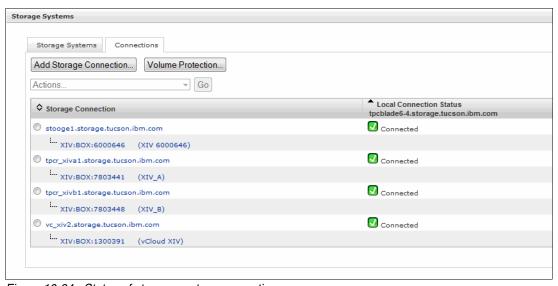


Figure 10-94 Status of storage system connections

If the connection was successful, the storage system is listed under the **Storage Systems** tab (Figure 10-95). It is also available for selection in the Add Copy Sets wizard for the sessions.

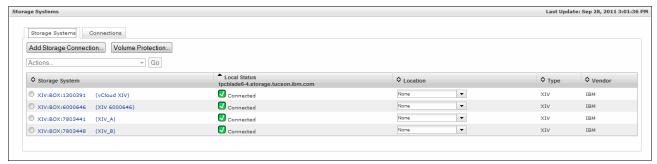


Figure 10-95 Successfully added XIV storage systems

Various panels within Tivoli Storage Productivity Center for Replication, such as those shown in Figure 10-94 and Figure 10-95, display the "Local Connection Status" for the added XIV storage systems. This status represents the status of the main connection between the Tivoli Storage Productivity Center for Replication server and the IP address that you added. It does not include the status of the additional IP connections to the XIV that are automatically found.

To view the status for all of the connections to an XIV storage system, select the radio button for the host name or IP address you added (the main connection), choose **View/Modify Connections Details** from the list of actions, and click **Go** (Figure 10-96). If you prefer, you can choose to click the link to the host name to go directly to the Connection Details panel (Figure 10-97) for a particular device instead.

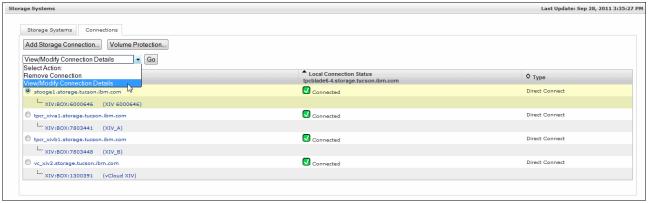


Figure 10-96 View/Modify Connections Details

Figure 10-97 shows the Local Connection Status for the main IP address you entered for the XIV storage system. It also lists Additional Module Connections on the right in the panel. The additional connections show the status of the connections to the other IP addresses that Tivoli Storage Productivity Center for Replication was able to find for the XIV storage system.

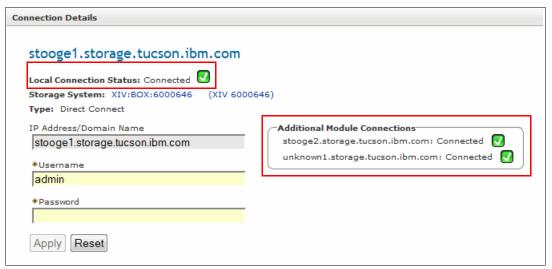


Figure 10-97 Connection Details for multiple IP addresses

Connections: The connection status values are all independent and do not roll up to provide an overall status value for the XIV storage system. The additional connections provide redundancy and failover for the nodes of the storage system.

Consistency groups

All three of the session types use XIV consistency groups. The consistency groups are created and named by Tivoli Storage Productivity Center for Replication. To determine the consistency group name for a session, we take the following approach:

- 1. Use the Tivoli Productivity Center for Replication session name (limited to 58 characters).
- 2. If that name already exists, use the session name with "_001" appended.
- 3. Keep trying " $_00x$ " names up to x=30.
- 4. If all of those names already exist, the consistency group will not be created.

Due to this naming convention, the consistency group names created might not be the same between XIV storage systems in a single session. You can have a consistency group named $mmSession_001$ on one XIV and a consistency group named $mmSession_002$ on the other. It depends on what consistency groups already exist on the individual XIV storage systems at the time Tivoli Storage Productivity Center for Replication attempts to create them.

The consistency group name is shown in the "Session Details" panel (Figure 10-98). It allows you to see what is being used on the XIV storage system, which can be important for debugging any issues.



Figure 10-98 Consistency groups listed in Session Details

You can also see the consistency group name in the Console log (Figure 10-99) as it is created.



Figure 10-99 Consistency group creation logged in the console

10.8.2 Snapshot sessions

Snapshot is a new session type that creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

The XIV storage system uses advanced snapshot architecture to create a large number of volume copies without affecting performance. By using the snapshot function to create a point-in-time copy, and to manage the copy, you can save storage. With the XIV storage system snapshots, no storage capacity is used by the snapshot until the source volume (or the snapshot) is changed.

Tip: The Snapshot session type is *only* available for XIV storage systems.

Configuration

XIV Snapshot session support is available with all Tivoli Storage Productivity Center editions. You must have the following environment to work with Snapshot sessions in Tivoli Storage Productivity Center for Replication:

- One or more XIV storage systems, with pools and volumes configured
- ► IP connectivity between the XIV storage system and the Tivoli Storage Productivity Center for Replication server

Limitations

The XIV Snapshot session limitations are as follows:

- Session name is limited to 58 characters.
- Consistency group is limited to 128 volumes.

This is not enforced by Tivoli Storage Productivity Center for Replication.

- All volumes from a session must be in the same pool.
- ▶ Volumes mapped to a host cannot be deleted while mapped.
- Locked volumes are read only.
- Snapshot groups can be automatically deleted.

This is based on deletion priority and pool space.

Data replication

After you have added connections to the XIV storage systems and met the configuration requirements listed above, you are ready to start creating sessions and adding copy sets to them. The general steps to set up a Snapshot session for data replication with Tivoli Storage Productivity Center for Replication are as follows:

- Create a Snapshot session: Select Sessions from the navigation tree and click Create Session to display the Create Session wizard. Choose XIV and Snapshot as you work through the wizard.
- Add copy sets on the XIV storage system to your session:
 On the Results page of the Create Session wizard, click Launch Add Copy Sets Wizard.
- Start the Snapshot session:
 In the Sessions panel, select the session name that you want to start, choose Create
 Snapshot from the actions list and click Go.

Detailed steps for creating a Snapshot and adding copy sets are available in Chapter 5, "Creating a snapshot session and adding copy sets" of the *Tivoli Storage Productivity Center Version 4.2.2 User's Guide*, SC27-2338.

Copy sets list

Role pair status is not provided for XIV system Snapshot sessions because the concept of role pairs does not exist for these sessions. Rather, a new (view) link has been provided to list the copy sets for a session. Figure 10-100 highlights where the new link is found in the Session Details.



Figure 10-100 View copy sets link

Click the (view) link to display a list of copy sets for the currently selected session, as shown in Figure 10-101.



Figure 10-101 List of copy sets for the session

New CLI commands for XIV Snapshot support

The following CLI commands are new. For more information about new and updated CLI commands for Tivoli Storage Productivity Center for Replication 4.2.2, see the IBM Tivoli Storage Productivity Center Command-line Interface Reference.

cmdsnapgrp	Use the cmdsnapgrp command to run a specific action against a snapshot	
	group that is in an XIV system Snapshot session. A snapshot group is a	
	grouping of snapshots of individual volumes in a consistency group at a	

specific point in time.

Ispool Use the **1spoo1** command to list all of the pools for an XIV storage

system.

Issnapgrp Use the **Issnapgrp** command to view snapshot groups that are in an XIV

system Snapshot session.

Issnapgrpactions Use the **Issnapgrpactions** command to specify the session and

snapshot group name for which you want to view available actions.

Issnapshots Use the **1ssnapshots** command to view snapshots that are in a snapshot

group in an XIV system session.

For XIV system sessions, you can provide the volume ID or the volume nickname as a parameter value when you add or remove copy sets by using the command line interface (CLI) commands mkcpset (Example 10-9) and rmcpset. In addition, you can include the XIV system volume ID or the volume nickname in a comma-separated value (CSV) file that is used to import copy set information.

You can import the CSV file by using the <code>importcsv</code> command or the Tivoli Storage Productivity Center for Replication graphical user interface (GUI). CSV files that are exported from Tivoli Storage Productivity Center for Replication for XIV system sessions include the volume nickname rather than the volume ID. CSV files are exported by using the <code>exportcsv</code> command.

Example 10-9 Adding a copy set using the nickname ID instead of the volume number ID

mkcpset -h1 XIV:VOL:7800000:db2 vol1 session

10.8.3 Metro Mirror Failover/Failback sessions

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

Metro Mirror replication maintains identical data in both the source and target. When a write is issued to the source copy, the changes made to the source data are propagated to the target before the write finishes posting. If the storage system goes down, Metro Mirror provides zero loss if data must be used from the recovery site.

If you are familiar with using the Metro Mirror session type with other supported storage systems, you will find the process within Tivoli Storage Productivity Center for Replication to be similar. We have highlighted areas that are unique to the XIV storage system.

Configuration

XIV Metro Mirror session support is available with Tivoli Storage Productivity Center Two Site Business Continuity. The license must be applied, but additional installation is not needed. You must have the following environment to work with Metro Mirror sessions:

- ► Two or more XIV storage systems
 - Pools and volumes configured
- ► IP connectivity between the XIV storage systems and the Tivoli Storage Productivity Center for Replication server
- Remote mirroring connectivity configured for the two XIV storage systems in the session
- Matching volumes on the source and target XIV storage systems
- All volumes in same pool on each host site

Reference: For information about XIV system configuration, see the *IBM XIV Storage System User Manual.* You can find this document and other XIV publications in the XIV Information Center. The user manual appears in "Publications" as "XCLI Reference."

http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp

If you have Tivoli Storage Productivity Center Standard Edition or Tivoli Storage Productivity Center Select, you can use the replication planning option of SAN Planner, available under Analytics in the navigation tree, to create the volume and session for use here. Details about SAN Planner support for XIV are covered in "SAN Planner with XIV replication" on page 494.

Limitations

The XIV Metro Mirror session limitations are as follows:

- ▶ Session name is limited to 58 characters.
- ► Consistency group is limited to 128 volumes.

This is not enforced by Tivoli Storage Productivity Center for Replication.

- ▶ All volumes from a session must be in the same pool.
- Volumes mapped to a host cannot be deleted while mapped.
- ► Locked volumes are read only.
- XIV hardware limited to 512 mirroring relationships.

Data replication

After you have added connections to the XIV storage systems and met the configuration requirements listed above, you are ready to start creating sessions and adding copy sets to them. The general steps to set up a Metro Mirror session for data replication with Tivoli Storage Productivity Center for Replication are as follows:

1. Create a Metro Mirror session.

Select **Sessions** from the navigation tree and click **Create Session** to display the Create Session wizard. Choose **XIV** and **Metro Mirror Failover/Failback** as you work through the wizard.

- Add copy sets on the XIV storage system to your session.
 On the Results page of the Create Session wizard, click Launch Add Copy Sets Wizard.
- 3. Start the Metro Mirror session.

In the Sessions panel, select the session name that you want to start, choose **Start H1->H2** from the actions list and click **Go**.

Tip: Metro Mirror sessions will assimilate existing hardware mirroring relationships, if the mirrored volumes are used as copy sets in the session.

More detailed instructions for creating a Metro Mirror session and adding copy sets are available in Chapter 5, "Creating a Metro Mirror session and adding copy sets" of the *Tivoli Storage Productivity Center Version 4.2.2 User's Guide*, SC27-2338-05.

10.8.4 Global Mirror Failover/Failback sessions

Global Mirror is a method of asynchronous, remote data replication that operates between two sites that are over 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

The data on the target is typically written a few seconds after the data is written to the source volumes. When a write is issued to the source copy, the change is propagated to the target copy, but subsequent changes are allowed to the source before the target verifies that it has received the change. Because consistent copies of data are formed on the secondary site at set intervals, data loss is determined by the amount of time since the last consistency group was formed. If your system stops, Global Mirror might lose some data that was being transmitted when the disaster occurred. Global Mirror still provides data consistency and data recoverability in the event of a disaster.

If you are familiar with using the Global Mirror session type with other supported storage systems, you will find the process within Tivoli Storage Productivity Center for Replication to be similar. We have highlighted areas that are unique to the XIV storage system.

Configuration

XIV Global Mirror session support is available with Tivoli Storage Productivity Center Two Site Business Continuity. The license must be applied, but additional installation is not needed. You must have the following environment to work with Global Mirror sessions:

- At least two XIV storage systems, with pools and volumes configured
- ► IP connectivity between the XIV storage systems and the Tivoli Storage Productivity Center for Replication server
- ► Remote mirroring connectivity configured for the two XIV storage systems in the session
- Matching volumes on the source and target XIV storage systems
- All volumes in same pool on same site

Reference: For information about XIV system configuration, see the *IBM XIV Storage System User Manual.* You can find this document and other XIV publications in the XIV InfoCenter. The user manual appears in "Publications" as "XCLI Reference":

http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp

If you have Tivoli Storage Productivity Center Standard Edition or Tivoli Storage Productivity Center Select, you can use the replication planning option of SAN Planner, available under Analytics in the navigation tree, to create the volume and session for use here. Details about SAN Planner support for XIV are covered in "SAN Planner with XIV replication" on page 494.

Limitations

The XIV Metro Mirror session limitations are as follows:

- Session name is limited to 58 characters.
- ► Consistency group is limited to 128 volumes.

This is not enforced by Tivoli Storage Productivity Center for Replication.

- ▶ All volumes from a session must be in the same pool.
- ▶ Volumes mapped to a host cannot be deleted while mapped.
- Locked volumes are read only.
- ▶ XIV hardware limited to 512 mirroring relationships.

Data replication

After you have added connections to the XIV storage systems and met the configuration requirements listed above, you are ready to start creating sessions and adding copy sets to them. The general steps to set up a Global Mirror session for data replication with Tivoli Storage Productivity Center for Replication are as follows:

1. Create a Global Mirror session.

Select **Sessions** from the navigation tree and click **Create Session** to display the Create Session wizard.

Much like the ESS/DS and SVC/Storwize V7000 storage systems, the XIV storage system has special Global Mirror options in Tivoli Storage Productivity Center for Replication. The Create Session wizard (Figure 10-102) presents these options after choosing **XIV** as the hardware type and **Global Mirror** as the session type.

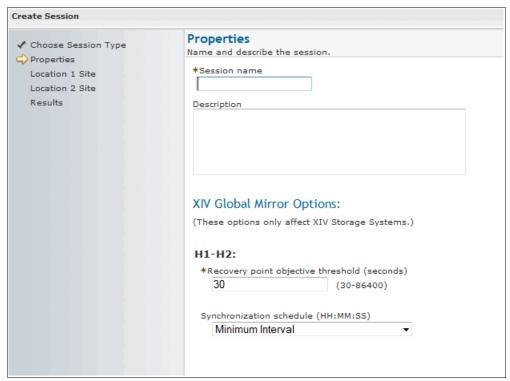


Figure 10-102 XIV Global Mirror options

Recovery point objective threshold (seconds):

Type the number of seconds that you want to set for the recovery point objective (RPO) threshold. RPO represents a measure of the maximum data loss that is acceptable in the event of a failure or unavailability of the master.

If the XIV system determines that the RPO is greater than this value, the session state becomes Severe. You can specify an RPO between 30 and 86400 seconds. The default is 30 seconds.

Synchronization schedule (HH:MM:SS):

Select an interval for the creation of an XIV system synchronization schedule. The XIV system attempts to form consistent points of data by taking automatic snapshots of the volumes in the session at this interval. The default is Minimum Interval, which is 20 seconds.

If you select Never, synchronization is not scheduled and the XIV system does not create consistency groups. When the XIV system determines that the RPO threshold has been passed, the session state becomes Severe.

2. Add copy sets on the XIV storage system to your session.

On the Results page of the Create Session wizard, click Launch Add Copy Sets Wizard.

3. Start the Global Mirror session.

In the Sessions panel, select the session name that you want to start, choose **Start H1->H2** from the actions list, and click **Go**.

Tip: Global Mirror sessions will assimilate existing hardware mirroring relationships, if the mirrored volumes are used as copy sets in the session.

More detailed instructions for creating a Global Mirror session and adding copy sets are available in the *Tivoli Storage Productivity Center Version 4.2.2 User's Guide*, SC27-2338-05. See Chapter 5 - "Creating a Global Mirror session and adding copy sets". Multiple Global Mirror sessions

Tivoli Storage Productivity Center for Replication 4.2 supports multiple Global Mirror sessions, allowing you to create more than one Global Mirror session per storage system. This allows you to fail over only data that is assigned to one host or application instead of forcing you to fail over all data if one host or application fails. This provides increased flexibility to control the scope of a failover operation and to assign separate options and attributes to each session.

Sessions: The multiple Global Mirror session limitations are as follows:

- Only one Global Mirror session is allowed per LSS.
- ► A maximum of 255 Global Mirror master sessions is allowed per box.
- ► A maximum of 255 Global Mirror subordinate sessions is allowed per box.

If you are planning to create multiple Global Mirror sessions, you can create them in the same way as with the previous Tivoli Storage Productivity Center for Replication releases. The difference with the multiple Global Mirror session is that when creating a session, it uses the first available Global Mirror session ID. The Global Mirror session number is chosen (in the background) by an algorithm that avoids already used Global Mirror session IDs. If the limit of Global Mirror sessions has been reached, a message is issued.

You can create Global Mirror sessions by using the Tivoli Storage Productivity Center for Replication GUI Create Session wizard or by using the mksess CLI command. The general steps to create a Global Mirror session and add copy sets are as follows:

- 1. Create Global Mirror session.
- 2. Add copy sets on an XIV storage system.
- 3. Start the Global Mirror session.

10.8.5 Troubleshooting

Even with careful planning and execution, you might still encounter errors when attempting these data replication tasks. This section provides guidance for some of the common errors that might occur.

Troubleshooting resources

The following files and tools can help you find additional information to help you dig into the errors:

► Log package (see 10.4, "Downloading log packages with the GUI" on page 384 for more details):

The log package does not require direct access to the Tivoli Storage Productivity Center for Replication file system. It contains logs with details regarding the actions in Tivoli Storage Productivity Center for Replication, such as *xivApiTrace*.

► Tivoli Storage Productivity Center for Replication Console:

The Console is a listing in the GUI of csmMessage.log that exists on the Tivoli Storage Productivity Center for Replication server. It can be opened by selecting **Console** from the navigation tree.

Figure 10-103 shows a sampling of the type of messages that are available in the Console. It can be used to identify steps that succeeded and allow you to isolate the step that failed. It also is a nice historical reference of actions against the Tivoli Storage Productivity Center for Replication server.



Figure 10-103 Sample Console listing

It can also be accessed using links provided during actions within the GUI. This can be handy for providing additional information at the time of the error. Click the **(Open Console)** link as shown in Figure 10-104 and Figure 10-105.



Figure 10-104 Console link for a running action



Figure 10-105 Console link for a completed action

Additionally, you can click the link to the message ID (for example, IWNR1026I) to bring up the message description.

Troubleshooting issues with Metro Mirror and Global Mirror sessions

This section addresses some of the issues you might encounter with Metro Mirror and Global Mirror sessions. Symptoms you might see, example errors, and steps to resolve are provided to aid you with troubleshooting.

Pool/snapshot size not large enough for global mirroring

► Symptoms (see Example 10-10 and Example 10-11):

Example 10-10 With volume IO pair errors after starting session, all pairs go suspended

IWNR2055W [Aug 31, 2011 9:16:45 AM] The pair in session volumespace for copy set XIV:VOL:7803441:100987 with source XIV:VOL:7803441:100987 (io_todd_3) and target XIV:VOL:7803448:101660(io_todd_3) in role pair H1-H2 was suspended due to a reason code of Master_Pool_Exhausted, but was not yet consistent; no action was taken on the session.

Example 10-11 Session after prepared moves severe/suspended on volume IO

IWNR2050E [Aug 26, 2011 9:48:42 AM] The pair in session testfullpool for copy set XIV:VOL:7803441:100985 with source XIV:VOL:7803441:100985(io_todd_1) and target XIV:VOL:7803448:101658(io_todd_1) in role pair H1-H2 was suspended unexpectedly with a reason code of Master Pool Exhausted.

Steps to resolve:

- Increase the size of the pool and snapshot space of the pool. The pool size must be more than three times the IO volumes total size for the pool. As long as there is enough hard space, the snapshot space is not as important.
- Refresh configuration for the XIV storage subsystem.
- Restart session.

Consistency group name already exists

Symptoms:

All operations are successful and the session runs correctly. There are no specific errors for this situation.

- Consistency group for the session will be listed as <session name> 001.
 - XIV automatically increments existing names.
- This is a cosmetic issue unless volumes intended for the session are also in the consistency group. (See the next troubleshooting topic.)

Steps to resolve:

- Terminate the session.
- Remove all of the volumes from the consistency group on the XIV.
- Remove the consistency group on the XIV.
- Restart the session.

Volumes already a part of consistency group

▶ Symptoms (see Example 10-12):

Example 10-12 Pair errors after starting session - VOLUME_BELONGS_TO_CG

IWNR2108E [Aug 26, 2011 12:17:42 PM] A hardware error occurred during the running of a command for the pair in session existingMM_CG for copy set XIV:VOL:7803441:100985 with source XIV:VOL:7803441:100985(io_todd_1) and target XIV:VOL:7803448:101658(io_todd_1) in role pair H1-H2. The hardware returned an error code of VOLUME_BELONGS_TO_CG.

- Steps to resolve:
 - Remove all of the volumes from the existing consistency group on the XIV.
 - Restart the session.

Volumes already a part of mirroring relationship

➤ Symptoms (see Example 10-13, Example 10-14, and Example 10-15):

Example 10-13 Volume is in relationship - VOLUME_IS_MASTER/SLAVE

IWNR2108E [Aug 26, 2011 1:25:56 PM] A hardware error occurred during the running of a command for the pair in session exisitingMirrors for copy set XIV:VOL:7803441:100985 with source XIV:VOL:7803441:100985(io_todd_1) and target XIV:VOL:7803448:101658(io_todd_1) in role pair H1-H2. The hardware returned an error code of VOLUME_IS_MASTER.

Example 10-14 Volume is in relationship - VOLUME_HAS_MIRROR

IWNR2108E [Aug 26, 2011 1:41:46 PM] A hardware error occurred during the running of a command for the pair in session exisitingMirrors for copy set XIV:VOL:7803441:100986 with source XIV:VOL:7803441:100986(io_todd_2) and target XIV:VOL:7803448:101659(io_todd_2) in role pair H1-H2. The hardware returned an error code of VOLUME_HAS_MIRROR.

Example 10-15 Volume mirror copy set is wrong copy type

IWNR2512E [Aug 26, 2011 1:25:56 PM] The volume mirror copy set XIV:VOL:7803441:100987 with source XIV:VOL:7803441:100987(io_todd_3) and target XIV:VOL:7803448:101660(io_todd_3) in session exisitingMirrors already existed on the hardware but was not the right copy type.

- Steps to resolve:
 - Deactivate the mirror.
 - Remove the mirror.
 - Restart the session.

Volumes have existing snapshots/replications

► Symptoms (see Example 10-16):

Example 10-16 Pair errors after starting the session.

IWNR2108E [Aug 30, 2011 11:07:15 AM] A hardware error occurred during the running of a command for the pair in session todd-mm for copy set XIV:VOL:7804988:580312 with source XIV:VOL:7804988:580312(io_todd_001) and target XIV:VOL:1302136:107903(io_todd_001) in role pair H1-H2. The hardware returned an error code of VOLUME_HAS_SNAPSHOTS.

- Steps to resolve:
 - Remove existing snapshots/replications on the XIV.
 - Restart the session.

Sessions go severe due to XIV hardware link errors

Symptoms (see Example 10-17 and Example 10-18):

Example 10-17 Prepared session and pairs go suspended or suspend after starting session

IWNR2061E [Sep 6, 2011 7:41:37 AM] The pair was suspended on the hardware because the source was disconnected from the target.

Example 10-18 Mirrors on XIV show RPO lagging

IWNR2750E [Sep 6, 2011 11:53:00 AM] Recovery Point Objective for session todd-gm-connectiondown has passed the threshold of 30 seconds.

- Steps to resolve:
 - Re-establish the link between the XIV storage systems.
 - Resolve lagging issues between the XIV storage systems.
 - Sessions will resolve on their own or you can restart the session.

Hardware troubleshooting: The following troubleshooting topics deal specifically with hardware configuration changes that might occur.

Changes to volume and pool name or size

- Symptoms:
 - Pair errors after starting the session.
 - Error returns BAD NAME.
- Steps to resolve:
 - Refresh the configuration for the XIV storage system.
 - Restart the session.

Changes to XIV hardware mirror relationships/consistency groups

Symptoms:

- Manually deactivate mirror → Tivoli Storage Productivity Center for Replication session is suspended.
- Remove mirror relationships → Tivoli Storage Productivity Center for Replication session is suspended.
- Remove volume from consistency group or group mirror → Tivoli Storage Productivity Center for Replication session will stay prepared, but it no longer has control of that volume for future commands.

Support: Tivoli Storage Productivity Center for Replication does not support handling any of these situations, but they are usually not fatal.

Steps to resolve:

 Each situation is unique, but in most cases restarting the session will resolve any manual manipulation of the hardware.

Changes to Global Mirror properties (RPO or Schedule)

Symptoms:

- There are no specific errors for this situation. You will see the old values for RPO
 Threshold and Sync Schedule when looking at the session properties or the hardware mirrors.
- Tivoli Storage Productivity Center for Replication will not automatically pick up changes made to Global Mirror properties on the hardware.

Steps to resolve:

- Make changes to the properties in the Tivoli Storage Productivity Center for Replication session panel.
- Restart the session.



XIV support

In this chapter, we describe the XIV device management and performance monitoring support provided by Tivoli Storage Productivity Center.

In Tivoli Storage Productivity Center V4.2, we use the XIV Native API and CLI to monitor and manage the XIV storage devices. This way provides additional resiliency in the communication between Tivoli Storage Productivity Center and the XIV devices compared to the SMI-S interface.

Tivoli Storage Productivity Center V4.2 adds support for performance management of IBM XIV Storage Systems.

For more details on various methods and functions that are available to monitor the IBM XIV Storage System, see the draft of *IBM XIV Storage System Architecture, Implementation, and Usage*, SG24-7659, which has an expected publish date of March 2012. The draft can be found at the following website:

http://www.ibm.com/redbooks

11.1 Supported firmware levels

The XIV firmware levels in Table 11-1 are supported by Tivoli Storage Productivity Center 4.2.

Table 11-1 Levels

XIV firmware version	Level of performance support
R10.0	Not supported for performance
R10.1	Limited support; can collect total I/0, total KB, and total time statistics
R10.2.2	Tracks statistics for read/write I/O, read/write KB, read/write time and read/write cache hits
R10.2.4	Enhanced volume-based and front-end performance metrics available with Tivoli Storage Productivity Center 4.2.1 fix pack 2 and higher

11.2 Adding XIV systems to Tivoli Storage Productivity Center

For details about how to add an XIV system to Tivoli Storage Productivity Center, see Chapter 7, "Device configuration: Introducing the Native API" on page 221.

Tip: Although an XIV system can have up to three IP addresses, only one needs to be added to Tivoli Storage Productivity Center. The other IP addresses are automatically discovered and will be used in a failover scenario.

11.3 XIV performance metrics

With each XIV firmware level supported by Tivoli Storage Productivity Center, the available performance metrics have increased. These statistics are uniform between the supported components unless specifically called out otherwise. The following tables separate the available performance metrics based on XIV versions.

Tip: You will have more performance metrics available when using the higher XIV versions. However, to use the additional metrics from XIV 10.2.4, you must be using Tivoli Storage Productivity Center 4.2.1 fix pack 2 or later.

The performance metrics listed in Table 11-2 are the limited metrics available for XIV 10.1 components. They apply to the volume, module, and subsystem components.

Table 11-2 Metrics for XIV 10.1

Performance metric	Description
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations for a component over a specified time interval
Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval

Performance metric	Description
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval
Overall Transfer Size	Average number of KB per I/O for read and write operations
Volume Utilization	The approximate utilization percentage of a volume over a specified time interval (the average percent of time that the volume was busy)

Tip: Performance statistics are only collected for the XIV interface modules, so there will be a maximum of 6 modules listed per device. If no volume I/O is routed through a particular module, there will not be any statistics available for that module.

The metrics listed in Table 11-3 are available for XIV 10.2.2 and higher for the volume, module, and subsystem components. This list includes the limited metrics that were available with XIV 10.1.

Table 11-3 Metrics for XIV 10.2.2 and higher

Performance metric	Description
Read/Write/Total I/O Rate(overall)	Average number of I/O operations per second for both sequential and nonsequential read and/or write operations for a component over a specified time interval
Read/Write/Total Cache Hit Percentage(overall)	Percentage of cache hits for both sequential and nonsequential read and/or write operations for a component over a specified time interval
Read/Write/Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read and/or write operations for a component over a specified time interval
Read/Write/Overall Response Time	Average number of milliseconds that it took to service each read and/or write operation for a component over a specified time interval
Read/Write/Overall Transfer Size	Average number of KB per I/O for read and/or write operations
Volume Utilization Percentage	The approximate utilization percentage of a volume computed over a specified time interval (the average percent of time that the volume was busy)

The metrics listed in Table 11-4 are available starting with XIV 10.2.4 and Tivoli Storage Productivity Center 4.2.1 fix pack 2. Unless marked as a port metric, these are available for the volume, module, and subsystem components.

Table 11-4 Additional metrics for XIV 10.2.4 and higher

Performance metric	Description
Small/Medium/Large/Very Large Transfers I/O Percentage	Percentage of I/O operations over a specified interval for particular data transfer sizes
Read/Write/Total Data Cache Hit Percentage	Percentage of read and/or write data that was read from or written to the cache over a specified time interval
Small/Medium/Large/Very Large Transfers Data Percentage	Percentage of data that was transferred over a specified interval for particular data transfer sizes
Read/Write/Overall Cache Hit Response Time	Average number of milliseconds that it takes to service each read and/or write cache hit operation over a specified time interval
Read/Write/Overall Cache Miss Response Time	Average number of milliseconds that it takes to service each read and/or write cache miss operation over a specified time interval
Small/Medium/Large/Very Large Transfers Response Time	Average number of milliseconds that it takes to service each I/O operation for various data transfer sizes
Port Send/Receive/Overall Response Time	Average number of milliseconds that it took to service each send (read) and/or receive (write) operation for a port over a specified time interval
Port Send/Receive/Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for send (read) and/or receive (write) operations for a port over a specified time interval
Port Send/Receive/Total I/O Rate	Average number of I/O operations per second for send and/or receive operations for a port over a specified time interval
Port Send/Receive/Overall Bandwidth Percentage	The approximate bandwidth utilization percentage computed for send and/or receive operations by a port based on its current negotiated speed

Tip: Unlike other storage subsystems, the XIV does not track performance of the underlying disk drives. Therefore, no back-end metrics are available in Tivoli Storage Productivity Center for XIV components.

For further information, see the following sections of *SAN Storage Performance Management Using Tivoli Storage Productivity Center*, SG24-7364:

- ► Appendix A. Rules of Thumb and suggested thresholds
- Appendix B. Performance Metrics and thresholds in Tivoli Storage Productivity Center performance reports

11.4 XIV storage provisioning

With Tivoli Storage Productivity Center, you can also provision storage on the XIV storage devices. This support includes creating, assigning, and deleting volumes.

To provision storage from an XIV system, navigate to **Disk Manager** \rightarrow **Storage Subsystems** and select the XIV device that you want to use (Figure 11-1).

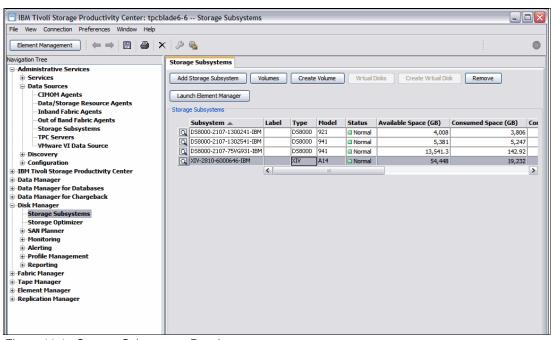


Figure 11-1 Storage Subsystems Panel

To create a new volume, click **Create Volume**. The Create Volume Wizard opens (Figure 11-2). Specify the number of volumes, size of each, and a volume name prefix.

RAID level: Ignore the RAID level selection for XIV, because the device does not use a traditional RAID scheme.

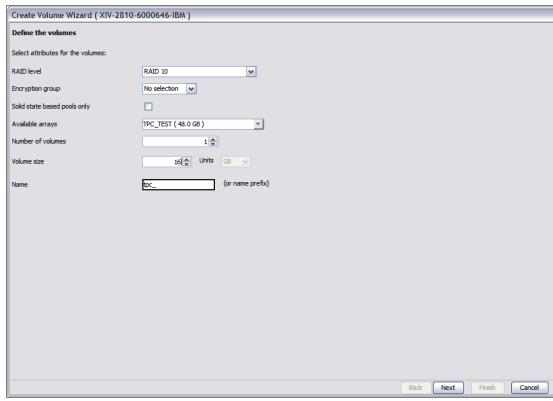


Figure 11-2 Create Volume Wizard

Tip: The Available arrays drop-down menu lists the XIV pools. New volumes are defined in the selected pool.

In the next panel (Figure 11-3), you assign the volumes to one or more host ports.

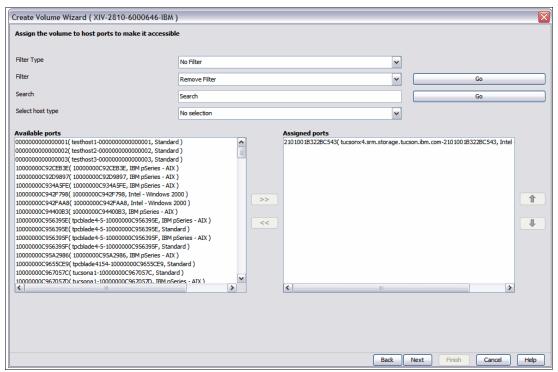


Figure 11-3 Create Volume Wizard host assignment

Attention: In this example, we do not use the Create Volume Wizard to make any SAN zoning changes.

The final panel in the Create Volume Wizard shows a summary of the proposed changes (Figure 11-4). Verify these. If they are correct, click **Finish**.

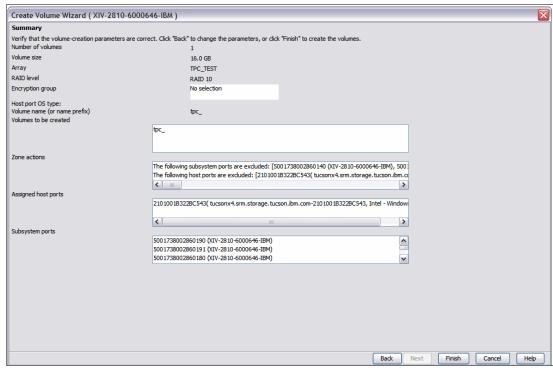


Figure 11-4 Create Volume Wizard Summary panel

A confirmation window opens (Figure 11-5). To see the status of the job, click **Yes** to open the Job Management panel.

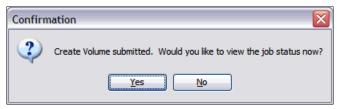


Figure 11-5 Create Volume Wizard Confirmation

Reference: SAN Planner can also be used to provision XIV storage. See Chapter 12, "SAN Planner" on page 437 for additional details.

11.5 XIV and SAN Planner

The SAN Planner is a function available with the Standard Edition license only. The following SAN Planner functions are available for XIV storage systems starting with Tivoli Storage Productivity Center 4.2.2:

- Performance considerations:
 - Workload profiles allow the flexibility to select volumes based on predefined performance requirements or custom performance requirements.
 - XIV performance monitors provide historical performance data needed to select the volumes meeting the performance requirements.
- Replication considerations:
 - Storage replication plans can be generated and implemented for all three replication types on XIV: Snapshot, Metro Mirror with Failover/Failback, and Global Mirror with Failover/Failback.

For further information about the SAN Planner functionality in Tivoli Storage Productivity Center, refer to Chapter 12, "SAN Planner" on page 437.

11.6 XIV and Storage Optimizer

The Storage Optimizer is a function available with the Standard Edition license only. It uses data in the Tivoli Storage Productivity Center database to analyze your storage subsystems to identify performance bottlenecks, and recommend changes to improve performance.

When Tivoli Storage Productivity Center 4.2 was released, it did not contain support for the Storage Optimizer function with XIV storage subsystems. Support was added later with the following versions:

- XIV 10.2.2 requires Tivoli Storage Productivity Center 4.2.1 fix pack 1 and higher.
- XIV 10.2.4 requires Tivoli Storage Productivity Center 4.2.1 fix pack 2 and higher.
- ► See the Supported Storage Products Matrix for additional details and future updates: https://www.ibm.com/support/docview.wss?uid=swg21386446

For further information about the Storage Optimizer functionality in Tivoli Storage Productivity Center, see Chapter 4, "Managing Storage Resources" in the IBM Tivoli Storage Productivity Center User's Guide, SC27-2338-05, and Chapter 7 of the IBM Tivoli Storage Productivity Center V4.1 Release Guide, SG24-7725.

Attention: Because XIV is not supported for the Storage Optimizer function with Tivoli Storage Productivity Center 4.1, the information covering this topic in the IBM Tivoli Storage Productivity Center V4.1 Release Guide will mention that XIV is not supported. Use this older guide for information regarding the concepts of Storage Optimizer, but refer to the Supported Storage Products Matrix for the updated list of supported devices.



SAN Planner

In this chapter, we discuss the function of SAN Planner in Tivoli Storage Productivity Center V4.2. With the SAN Planner, you can do end-to-end planning of fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, zone sets, storage resource groups (SRGs), and replication. After a plan is made, the client has the choice of having it implemented by SAN Planner.

SAN Planner supports TotalStorage Enterprise Storage Server, IBM System Storage DS6000, IBM System Storage DS8000, IBM System Storage SAN Volume Controller, IBM Storwize V7000, and IBM XIV Storage System.

SAN Planner supports the *Space Only* workload profile option for any other storage system supported by Tivoli Storage Productivity Center.

Terminology used: Certain system menus in this chapter issue "recommendations" that are meant to serve as guidelines for your choices.

12.1 Purpose of SAN Planner

When you want to provision storage to the host, you can use SAN Planner to provide its best recommendations for creating volumes in the storage subsystems. You can also do multipath planning and zoning using SAN Planner. In addition, if you want to replicate the volumes, you can use SAN Planner to provide its recommendation for that too.

After you have the SAN Planner's recommendation, also called planner output, you can review the plan and choose to execute it, in which case the planner will create a job to make the changes in the environment based on the plan output. Alternatively, you can vary the input provided to the SAN Planner to get multiple possible plans.

12.2 New features and functions

Tivoli Storage Productivity Center V4.2 has a new SAN Planner wizard, which has been enhanced to provide the following functions:

► SAN Volume Controller (SVC) and Storwize V7000 support with provisioning and workload profiles:

The SAN Planner recommendations are limited to SAN Volume Controller and Storwize V7000 front-end operations only; it does not support back-end operations such as the creation of new MDisks or the creation of MDisk groups. The support includes the creation and provisioning of VDisks with the recommended I/O group and preferred node for each VDisk.

► Space Only planning:

This function is for all storage subsystems supported by Tivoli Storage Productivity Center.

► Disaster Recovery Planner:

This function is for devices supported by Tivoli Storage Productivity Center for Replication (DS8000, DS6000, Enterprise Storage Server, SAN Volume Controller, Storwize V7000 and XIV).

► Space-efficient Volumes:

The SAN Planner now has an option to provision space-efficient volumes on supported storage subsystems (SAN Volume Controller, Storwize V7000, XIV, and DS8000).

Encrypted Volumes:

The SAN Planner now supports the creation or assignment of encrypted volumes for the DS8000 and SAN Volume Controller or Storwize V7000 (if a DS8000 is used as the back-end device).

► Candidate Storage Resource Group:

The SAN Planner now supports candidate storage resource groups, which are containers of storage resources. When you provide a candidate storage group as input for the planner, it uses the storage subsystems, pools, and volumes from that storage resource group for provisioning new storage.

12.3 Prerequisites for using SAN Planner

The prerequisites for using the SAN Planner are as follows:

- Discover and probe the storage subsystems in Tivoli Storage Productivity Center.
- ► Start performance monitors on the subsystems, which the SAN Planner supports for performance based planning.
- ▶ Discover and probe the fabrics in Tivoli Storage Productivity Center.
- ▶ Install Storage Resource agent (or legacy Data agent) on the host systems.
- ► Make sure that the supported IBM SDD driver (or Linux DM-Multipath driver or HP-UX native multipathing feature) is installed on the host systems (for multi-path planning).
- ► Add the storage subsystems to Tivoli Storage Productivity Center for Replication, for replication planning.
- ▶ Make sure proper connectivity exists between the subsystems for replication.
- ► For VDisk mirroring planning, make sure you have IO Group Memory configured on the SVC or Storwize V7000.
- ► For performance-based virtual disk planning, make sure you have performance data collected for all the back-end subsystems of SVC or Storwize V7000. In addition, the back-end subsystems should be one or more of DS8000, DS6000, DS4000 or Enterprise Storage Server only.
- ► For replication planning, make sure you have an appropriate Tivoli Storage Productivity Center for Replication license installed and Subsystem Device Feature codes enabled.
- ► For replication planning with XIV, make sure to have XIV System Software Version 10.2.4.b or higher, and Tivoli Storage Productivity Center version V4.2.2 or higher.

12.4 Supported storage subsystems in SAN Planner

If you want to only do storage provisioning without any replication or performance considerations, any storage subsystem managed by Tivoli Storage Productivity Center is supported.

12.4.1 Performance-based planning

As of Tivoli Storage Productivity Center V4.2, SAN Planner supports only IBM DS8000, DS6000, Enterprise Storage Server, SVC, Storwize V7000 and XIV for performance-based planning.

12.4.2 Space Only planning

If you have any other storage subsystem to use in planner that is supported by Tivoli Storage Productivity Center but is not listed above, select the **Space_Only** workload profile. The SAN Planner will recommend the new volumes solely based on space availability of the pools in the subsystem. Of course, you can also use the Space_Only profile for these subsystems for which performance-based planning is supported.

12.4.3 Replication planning

If you want to do replication planning also, the supported subsystems are IBM DS8000, DS6000, Enterprise Storage Server, SVC, Storwize V7000, and XIV.

12.5 Storage Resource Groups Monitoring and Alerting

To use all of the SAN Planner features, be sure to use Storage Resource Groups to group your storage entities.

The Storage Resource Groups concept offers a way to logically group multiple storage related entities in a single object. The Storage Resource Groups can also be used for monitoring purposes. This activity is mainly carried out from the Topology View. L0 and L2 Storage Resource Group views have been added to the Topology View. Figure 12-1 shows the Storage Resource Group L2 view.

Notice the icon that is related to the Storage Resource Group operational status. This information is "propagated" upwards and "health" of the group members becomes the "health" of the group itself. This information is helpful to quickly check the status of a group if, for example, the group logically represents an application.

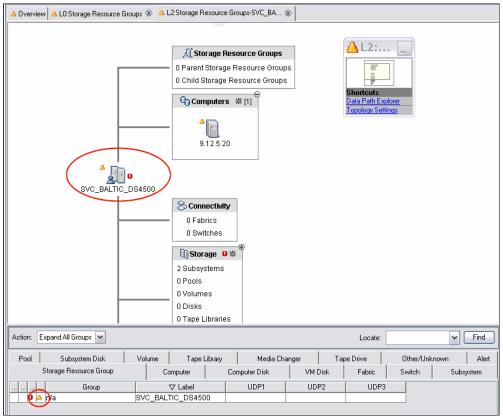


Figure 12-1 L2 Topology View for Storage Resource Groups

The information about the alerts is also propagated from the members upwards to the group. The corresponding graphical indicator is visible next to the Storage Resource Group icon and an alert overlay is made available, as shown in Figure 12-2.

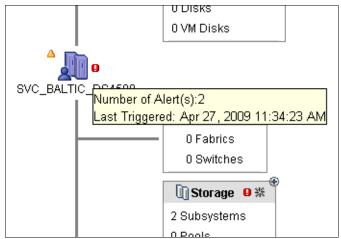


Figure 12-2 Alert Overlay for Storage Resource Group

Additionally the Storage Resource Groups health status is reported in the Storage Resource Group Management panel, as shown in Figure 12-3.

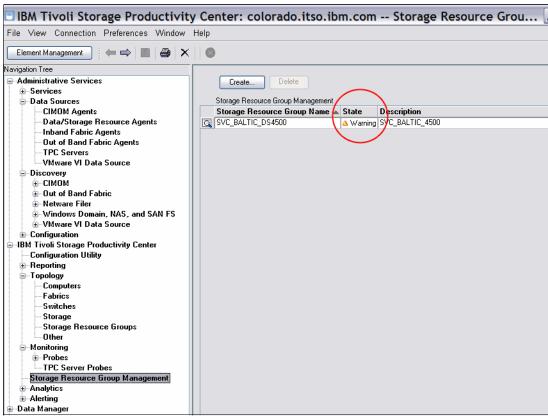


Figure 12-3 Storage Resource Group Management panel

12.6 Creating a Space Only SAN Planner recommendation

To create a new SAN plan, complete the following steps:

 Navigate to IBM Tivoli Storage Productivity Center → Analytics → SAN Planner, right-click and select Create Plan (Figure 12-4).

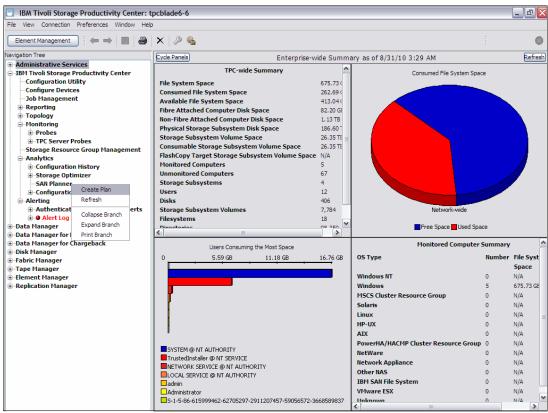


Figure 12-4 Create a new plan

2. The SAN Planner Wizard opens (Figure 12-5). Click Next.

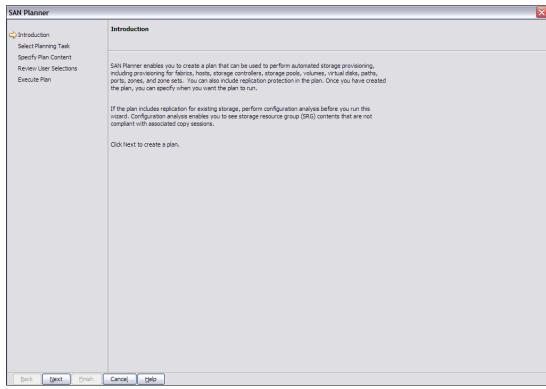


Figure 12-5 SAN Planner Introduction Panel

3. Select the type of planning you want to perform (Figure 12-6).

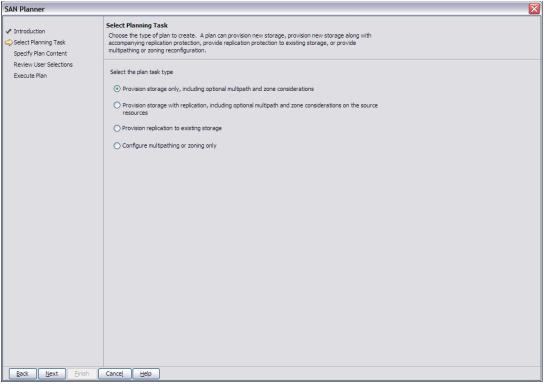


Figure 12-6 SAN Planner Select Planning Task

- Here, you can choose to provision storage only, provision storage with replication, setup replication on existing storage or configure multipathing or zoning to existing storage.
- 4. In this example, we provision storage only. Click **Next**. You have the option to select what type of new storage to provision, and whether to include multipathing or zoning considerations in the plan (Figure 12-7). Select either **Virtual Disks** to provision SAN Volume Controller or Storwize V7000 storage, or **Volumes** for any other type of storage device.

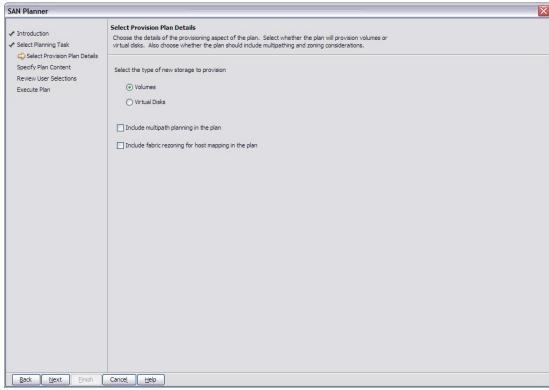


Figure 12-7 Select Provision Plan Details Panel

5. In this example, we provision volumes but did not include the multipath and zoning considerations. Click **Next**. You then have the option to specify the planner content (Figure 12-8).

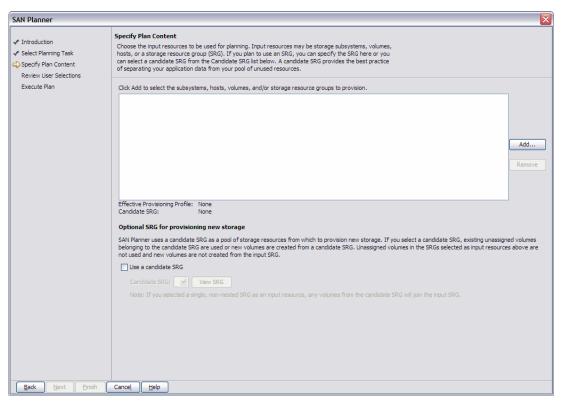


Figure 12-8 Specify Plan Content Panel

6. Click **Add** to add storage devices and hosts to provision storage on. The Topology Viewer opens (Figure 12-9).

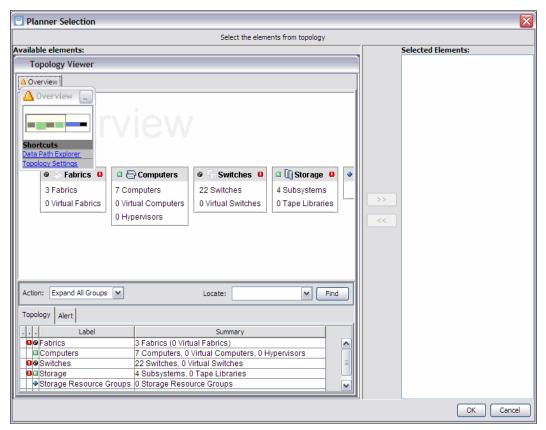


Figure 12-9 Topology Viewer selection for SAN Planner

- 7. In this example, we choose two DS8000s to allocate storage from and a single Windows host on which to assign the storage:
 - a. Select each storage device and click the >> button to move it to the selected elements pane (Figure 12-10).

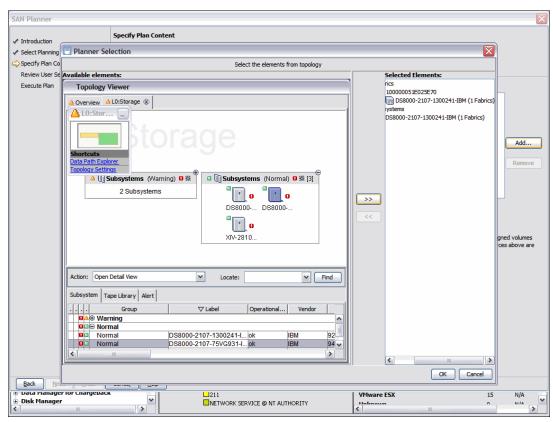


Figure 12-10 SAN Planner Storage Selection

b. Select the host and move it to the selected elements pane (Figure 12-11).

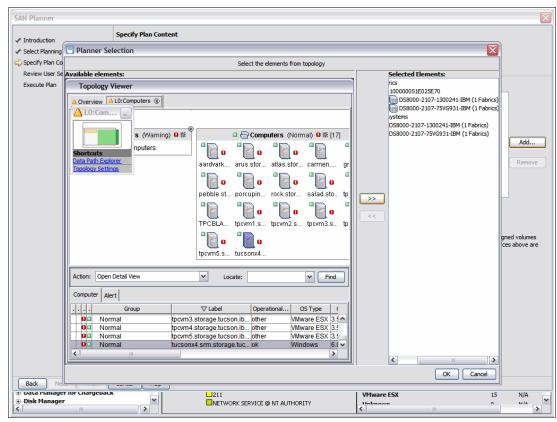


Figure 12-11 SAN Planner Host Selection

8. After you have selected the appropriate components, click **OK**. The Specify Plan Content window is updated with the entities you have selected (Figure 12-12).

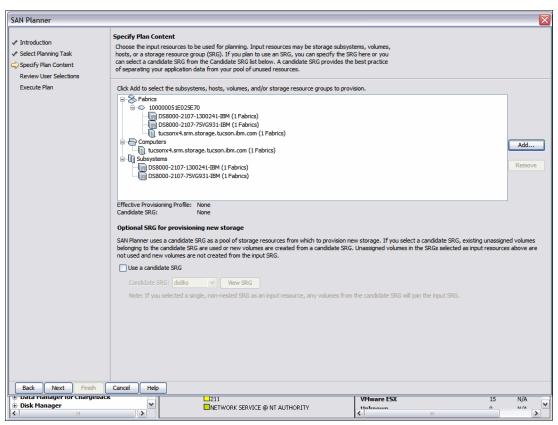


Figure 12-12 Populated SAN Plan Content Panel

 After you verify the plan contents, click Next. The capacity plan panel opens (Figure 12-13). Here you can set the total storage capacity that you want allocated, and how you want that storage carved into multiple volumes.

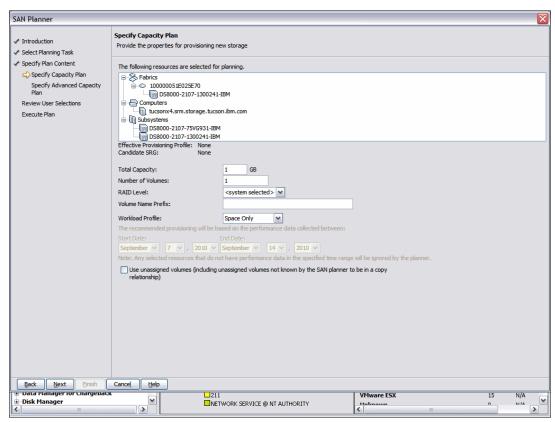


Figure 12-13 Specify Capacity Plan panel

For this example, we allocate 10 GB of storage, divided into five volumes. We let the system choose the RAID level, and select the default Space Only workload profile (Figure 12-14).

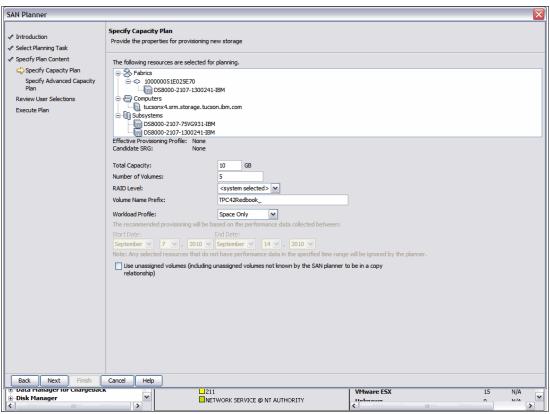


Figure 12-14 Specify Plan Content panel

10. Click **Next** on the wizard panel. The Advanced Capacity Plan panel opens. Here, you can select to use thin provisioning, solid state disks, or disk encryption. For this example, we do not select any advanced options (Figure 12-15).

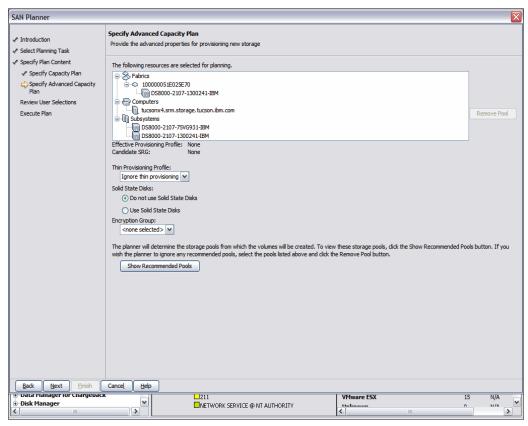


Figure 12-15 Advanced Capacity Plan panel

11. Click **Next**. A confirmation panel opens and you can validate the resources and options selected (Figure 12-16).

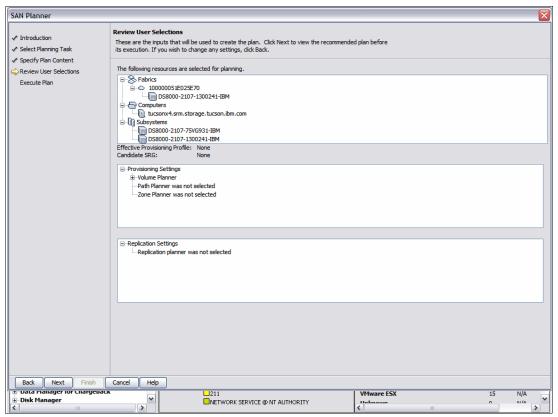


Figure 12-16 Review User Selection panel

12. Click **Next** in this panel. A recommendation that is based on your inputs is generated (Figure 12-17).

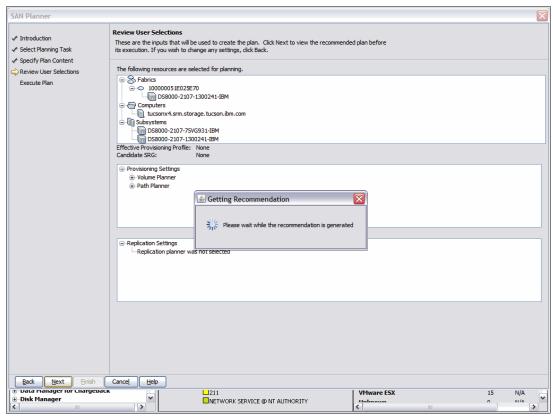


Figure 12-17 Generating Recommendation

After the recommendation is generated, a list of proposed changes is displayed (Figure 12-18). These changes include volume creations, volume assignments, and zoning changes. You can validate that all these changes are expected and if you want to change any of your inputs, click the **back** button in the wizard.

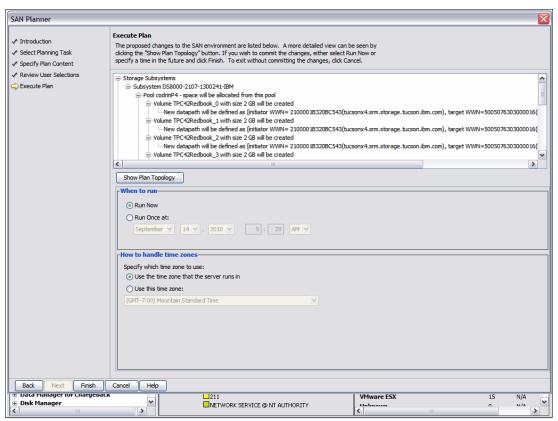


Figure 12-18 Execute Plan Panel

In the previous panel (Figure 12-18), you have the option to either run the plan now, or schedule it to run at a future time (for example, during your next scheduled change window). In this example, we choose to run it now. Click **Finish**. The plan is executed, and the volumes are created and assigned.

13. The wizard prompts you for a plan name (Figure 12-19).

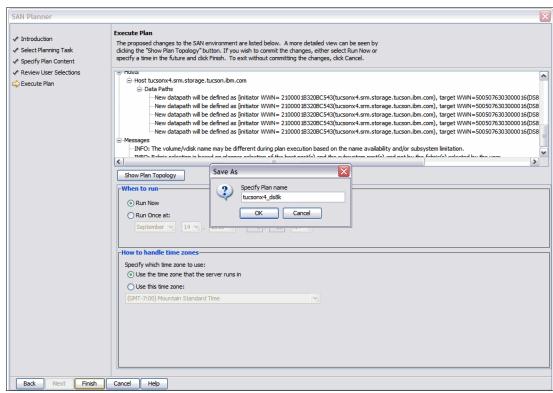


Figure 12-19 Specify Plan Name

14. After you save the plan, you can view the submitted job status through the Job Management panel (Figure 12-20).

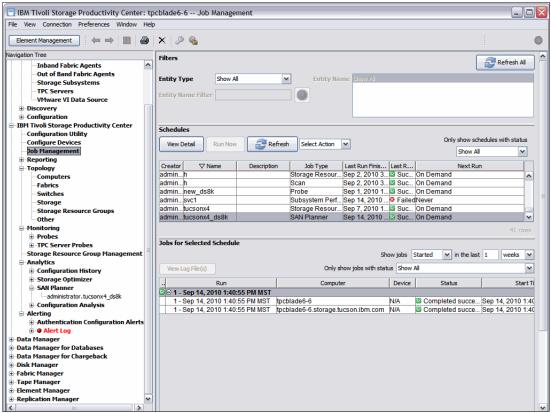


Figure 12-20 Job Management panel

12.7 Creating a DR Planner recommendation

In this section, we describe how to use the Disaster Recovery (DR) Planner to provision storage with replication.

12.7.1 Introduction

For replication planning, you must add all the subsystems, which are to be used in planning, to Tivoli Storage Productivity Center for Replication. You must also ensure that connectivity is established between the subsystems.

In addition to provisioning storage with multipathing and zoning options, when you select replication options in the wizard, SAN Planner provisions replication for your storage resources. Resilient resource planning is available only for devices that are supported by Tivoli Storage Productivity Center for Replication.

Depending on the storage resources you are planning for, you can select from the following copy services session types:

- FlashCopy
- Snapshot
- VDisk Mirroring

- Synchronous Metro Mirror failover/failback
- ► Synchronous Metro Mirror failover/failback with practice
- Synchronous Metro Mirror single direction
- ► Asynchronous Global Mirror either direction with two-site practice
- Asynchronous Global Mirror failover/failback
- Asynchronous Global Mirror failover/failback with practice
- ► Asynchronous Global Mirror single direction
- ► Three-site Metro Global Mirror with practice
- ► Three-site Metro Global Mirror

Replication can also be set up to existing storage. This option allows you to extend existing replication sessions by adding more storage resources and protecting it. The new volumes are added to existing replication sessions. SAN Planner ensures that source volumes are added to the source side of a copy relationship and target volumes are added to the target side.

12.7.2 Disaster Recovery Planner scenario

The following example is a DR Planner scenario, using Synchronous Metro Mirror failover/failback:

 Add the storage devices to a candidate Storage Resource Group. Under IBM Tivoli Storage Productivity Center, select Storage Resource Group Management, and click Create (Figure 12-21).

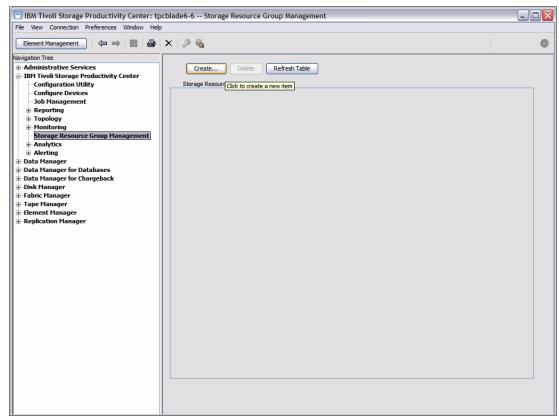


Figure 12-21 Storage Resource Group

The Create Storage Resource Group window opens (Figure 12-22).

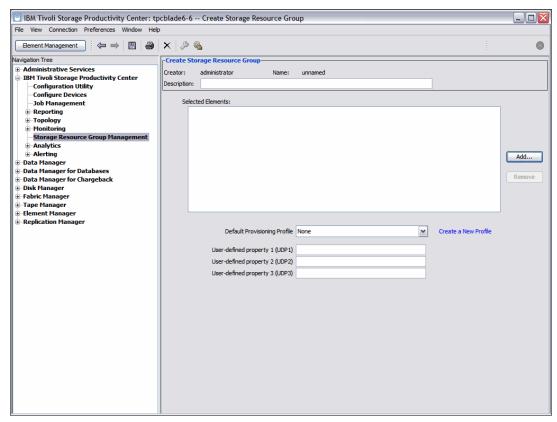


Figure 12-22 Create Storage Resource Group panel

Figure 12-23 shows the following information and objects:

- Creator:

Displays the user name of the creator.

- Name:

Displays the name of the storage resource group or unnamed, if it is not yet named.

- Description:

Optional: Displays the user defined description for the storage resource group.

- Selected Elements:

Lists the elements selected to be members of this storage resource group.

– Add:

Adds one or more selected elements to the list. The Storage resource group element selection panel is displayed.

- Remove:

Removes one or more selected elements from the list.

- Default Provisioning Profile:

Lists the available provisioning profiles that can be associated with storage resource groups. The list also includes "None." If this storage resource group is used as input to the SAN Planner, the settings defined in this profile will be used to pre-populate the planner inputs.

- Create a New Profile:

Launches the Provisioning Profile creation wizard. When you complete the wizard, the Provisioning Profile list is updated.

User defined property 1 (UDP1):

Specifies any user-defined properties that will be used by the Topology Viewer to provide custom groupings.

User defined property 2 (UDP2):

Specifies any user-defined properties that will be used by the Topology Viewer to provide custom groupings.

- User defined property 3 (UDP3):

Specifies any user-defined properties that will be used by the Topology Viewer to provide custom groupings.

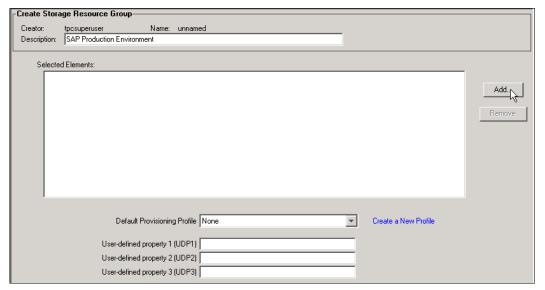


Figure 12-23 Create Storage Resource Group

In this example, we add two DS8000s to the candidate group (Figure 12-24).

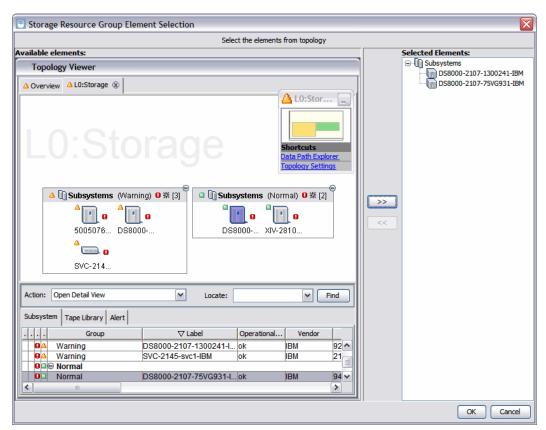


Figure 12-24 Topology Selector for new Storage Resource Group

2. We also create a new provisioning profile to be used with this SRG. Click **Create a New Profile** to open the Create Provisioning Profile window (Figure 12-25).



Figure 12-25 Create Provisioning Profile Panel

3. In this example, we create a new profile without using an existing profile. Keep the default options and click **Next**. The next panel asks for the volume settings. We provision 10 GB of storage, divided into two volumes. All other options remain as default values (Figure 12-26).

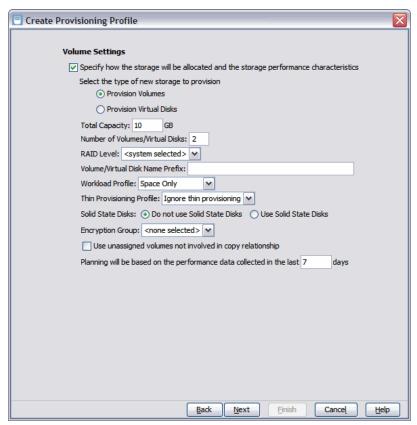


Figure 12-26 Create Provisioning Profile Volume Settings

4. We do not do any multipath modification, so we clear the **Setup Multipath options** check box and click **Next** (Figure 12-27).



Figure 12-27 Create Provisioning Profile Multipath Settings

5. For this example, we also do not do any zoning modifications, so we clear the zoning option (Figure 12-28).



Figure 12-28 Provisioning Profile Zoning/Security Settings

6. Click **Finish** to create the provisioning profile. You are again back at the Create Storage Resource Group panel. Make sure to select the provisioning panel that you just created under the Default Provisioning Profile (Figure 12-29).

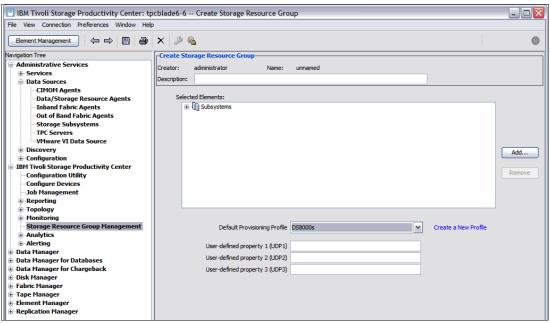


Figure 12-29 Create Storage Resource Group panel

7. Save the Storage Resource Group by either clicking the disk icon or **File** \rightarrow **Save** (Figure 12-30).

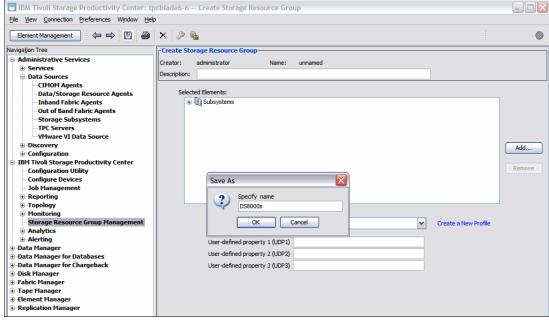


Figure 12-30 Save Storage Resource Group

8. After you have created the Storage Resource Group, you can start the SAN Planner by expanding IBM Tivoli Storage Productivity Center → Analytics, right-clicking SAN Planner, and selecting Create Plan (Figure 12-31).

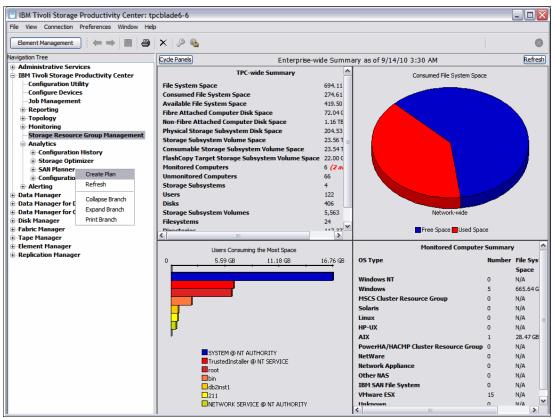


Figure 12-31 Create new SAN Planner

9. For the plan task type, select **Provision storage with replication** on the Select Planning Task panel and click **Next** (Figure 12-32).

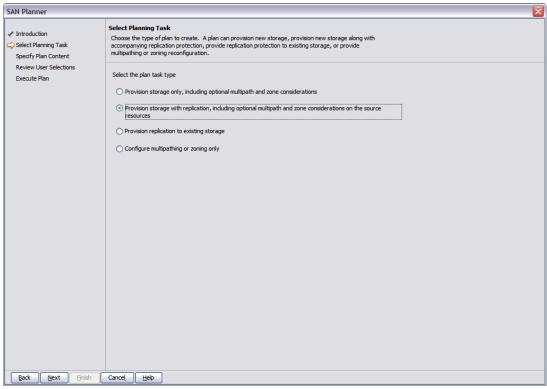


Figure 12-32 Select SAN Planning Task

10.On the Select Provision Plan Details window, we choose to create new volumes, without multipath or zoning considerations (Figure 12-33).

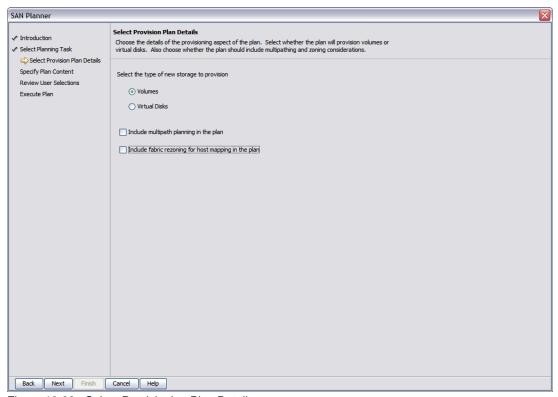


Figure 12-33 Select Provisioning Plan Details

11. For the plan content, we choose a candidate SRG instead of manually selecting entities for storage to be provisioned from (Figure 12-34).

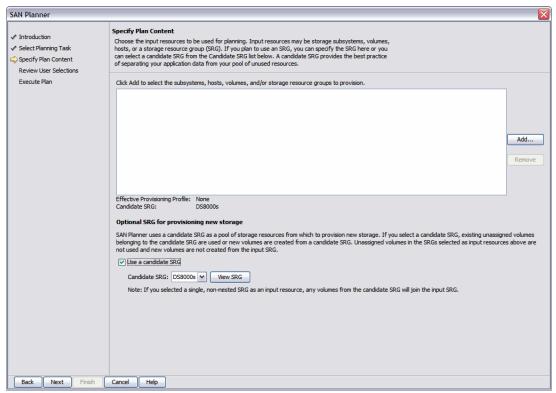


Figure 12-34 Specify Plan Content

12. We choose to provision 10 GB of capacity between two volumes and select Space Only workload profile because in this example we are considering capacity only (Figure 12-35).

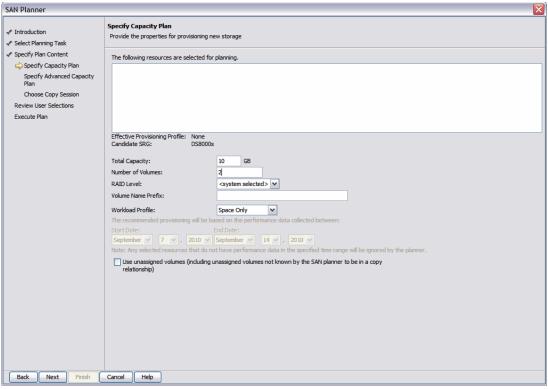


Figure 12-35 Specify Capacity Plan

13.In this example, we do not select any advanced plan options (Figure 12-36). Leave the defaults and click **Next**.

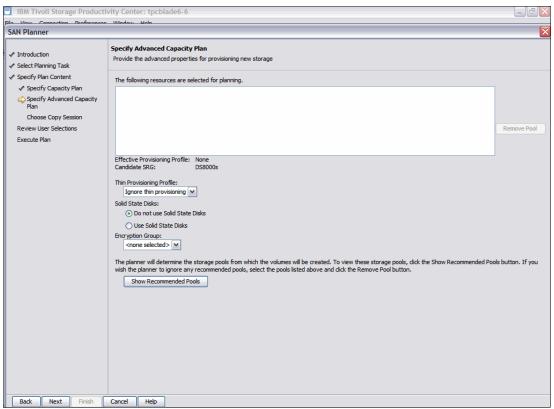


Figure 12-36 Specify Advanced Capacity Plan

14. The next panel shows any existing replication sessions, and if applicable, allows you to select to use an existing replication session. For this example, we create a new replication session, which is the default (Figure 12-37). Click **Next**.

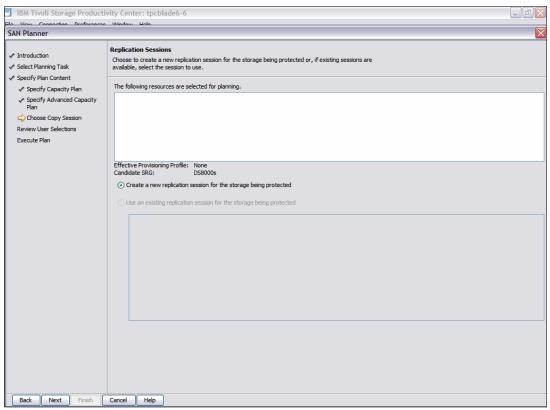


Figure 12-37 Replication Sessions Panel

15. The session properties panel opens. You can input the replication session name and choose the session type. Only the session types supported by the selected subsystems will appear in the list. These session properties are what you will see in Tivoli Storage Productivity Center for Replication (Figure 12-38).

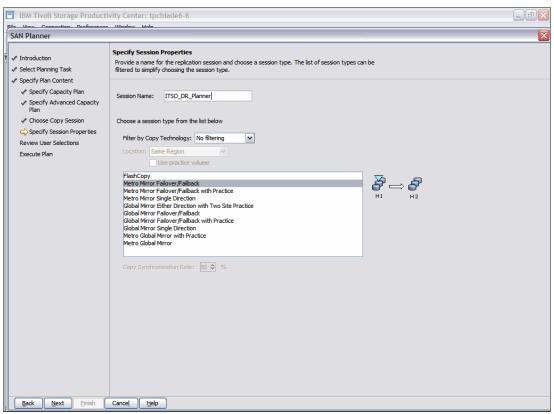


Figure 12-38 Specify Session Properties panel

16. The next panel requests that you chose the secondary location. In this case, we choose a secondary candidate storage resource group. This is where the target replication storage will be allocated from (Figure 12-39).

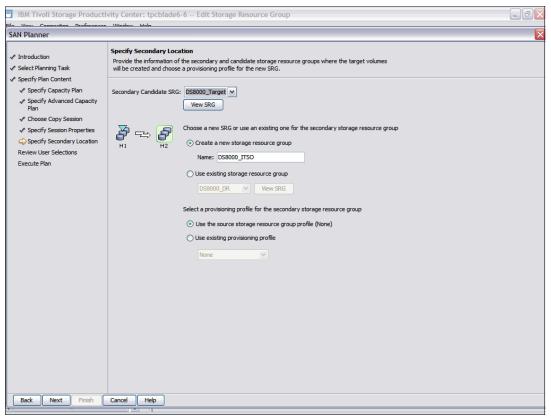


Figure 12-39 Specify Secondary (Target) Location panel

17. Use the Review User Selections panel to validate the wizard inputs before the recommendation is generated (Figure 12-40).

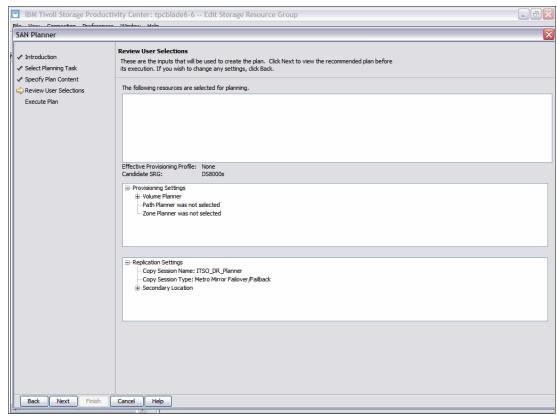


Figure 12-40 Review User Selections panel

18. Validate that these inputs are expected, and then click Next to generate the recommendation (Figure 12-41).

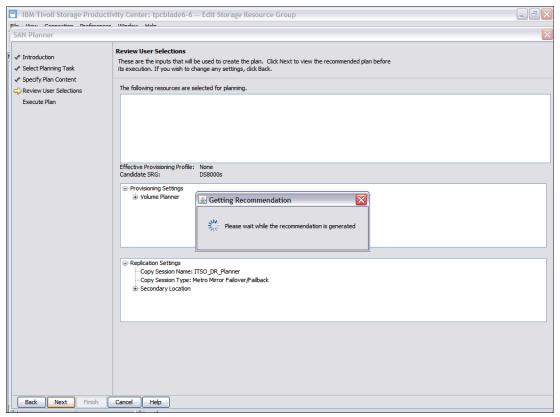


Figure 12-41 Generating Recommendation panel

19. The final panel within the SAN Planner wizard lists all recommended changes (Figure 12-42). Within this panel, you can also view the proposed changes in a graphical view by clicking **Show Plan Topology**.

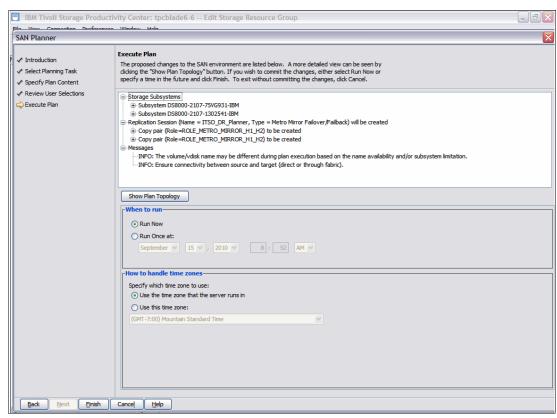


Figure 12-42 Execute Plan panel

🖯 IBM Tivoli Storage Productivity Center: tpcblade6-6 -- Edit Storage Resource Group Show Plan Topology Plan Preview 🛞 **⇔**Exe **⊟** Computers Subsystems # [2] 0 Computers 0 Fabrics Volumes (FB_P2) ∰ 1 Volume EPCPorts (All) ♣ Pools (All) ※ 0 FCPorts 2 Pools **\$\text{\$\psi\$}\text{Volumes}\text{ (FB_P1) } \text{\psi}\$** 1 Volume Action: Expand All Groups Find Locate: Computer | HBA | Computer Disk | FCPort | Fabric | Switch | Subsystem | Subsystem Disk | Pool | Volume | Replication Volumes | Alert | Data Path Operational... OS Type OS Version Hostname IP Address Computer T... Hypervisor

The SAN Planner output is displayed, but not the current environment (Figure 12-43).

Figure 12-43 SAN Planner Topology view

- 20. After you have validated the recommendation, you are able to either run the job immediately or schedule it to run in the future. You then click **Finish** and the job is saved. If you selected run immediately, the job is saved and started.
 - You can check on the status of the SAN Planner job within the Job Management panel (see Figure 12-44).

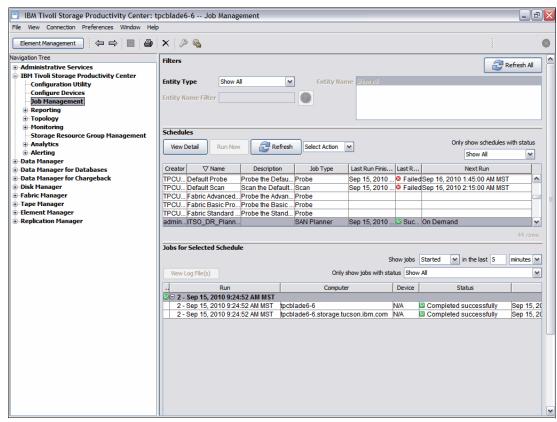


Figure 12-44 Job Management view for SAN Planner job

You can see that the job completed successfully and that the replication relationship was created in Tivoli Storage Productivity Center for Replication (Example 12-1).

Example 12-1 SAN Planner Execution Job

```
9/15/10 9:24:52 AM STS0306I: Job queued for processing. Waiting for idle thread.
9/15/10 9:24:52 AM HWNLM0001I: HWNLM0001I An integrated SAN Planner job started with
schedule administrator. ITSO DR Planner
9/15/10 9:24:52 AM HWNLM0011I: HWNLM0011I Started to create storage volumes.
HWN020001I Operation createStorageVolumes processed successfully.
HWNEP0138I External process was successfully executed for device 2107.75VG931.
HWNEP0138I External process was successfully executed for device 2107.75VG931.
HWNEP0138I External process was successfully executed for device 2107.1302541.
HWNEP0138I External process was successfully executed for device 2107.1302541.
9/15/10 9:25:31 AM HWNLM0013I: HWNLM0013I Completed creating storage volumes.
9/15/10 9:25:33 AM HWNLM810I: HWNLM810I Storage Subsystem Configuration refreshed
successfully in Replication Manager Subsystem=DS8000-941-75VG931.
9/15/10 9:25:42 AM HWNLM810I: HWNLM810I Storage Subsystem Configuration refreshed
successfully in Replication Manager Subsystem=DS8000-941-1302541.
9/15/10 9:25:42 AM HWNLM803I: HWNLM803I Replication Session was created successfully
ITSO DR Planner.
9/15/10 9:25:43 AM HWNLM807I: HWNLM807I CopySets Added to Session successfully
ITSO DR Planner.
9/15/10 9:25:43 AM HWNLM813I: HWNLM813I Replication Session was started successfully
ITSO DR Planner.
9/15/10 9:25:43 AM HWNLM0003I: HWNLM0003I The integrated SAN Planner job completed.
```

You can also go to the Tivoli Storage Productivity Center for Replication Console and see that the session was created and successfully started (Figure 12-45).

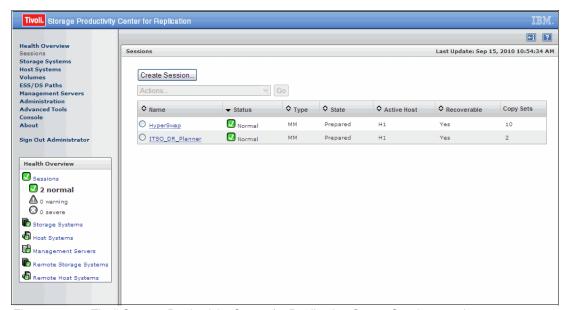


Figure 12-45 Tivoli Storage Productivity Center for Replication Create Session panel

12.8 SAN Planner with SVC and Storwize V7000

In Tivoli Storage Productivity Center V4.2, the SAN Planner supports IBM SAN Volume Controller and Storwize V7000. They are supported with both the workload profiles and in Space Only modes.

In this example, we show you how to allocate storage from a SAN Volume Controller, and set up VDisk Mirroring.

Licensing: VDisk Mirroring is done natively on the IBM SAN Volume Controller and Storwize V7000, and therefore, a Tivoli Storage Productivity Center for Replication license is not required.

Complete the following steps:

 To start the SAN Planner task, expand IBM Tivoli Storage Productivity Center → Analytics. Right-click SAN Planner, and select Create Plan (Figure 12-46).

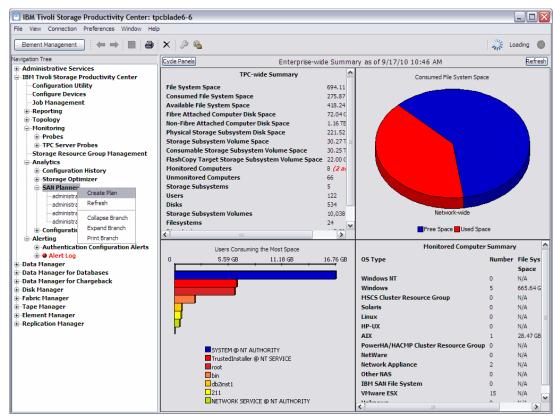


Figure 12-46 Create SAN Planner job

2. In the Introduction page, click **Next** to open the Select Planning Task window. Here, select **Provision storage with replication** (Figure 12-47).

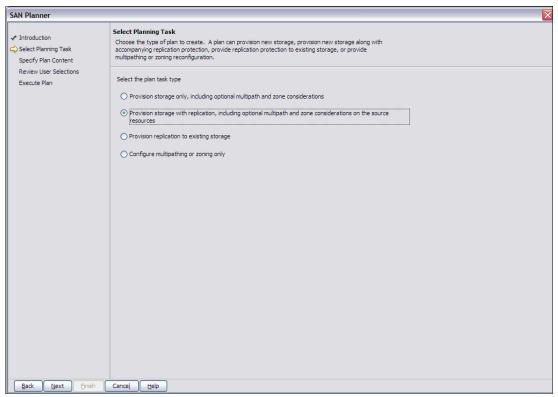


Figure 12-47 Select Planning Task

3. In the next panel, we choose to provision virtual disks (select **Virtual Disks**), and clear the check boxes for the multipath planning and zone (Figure 12-48).

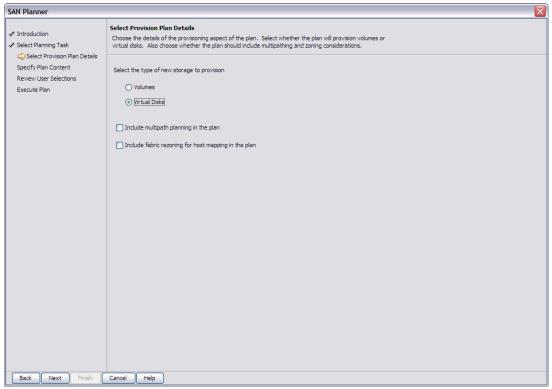


Figure 12-48 Select Provision Plan Details

4. We choose a single SAN Volume Controller to allocate storage from on the Specify Plan Content panel (Figure 12-49).

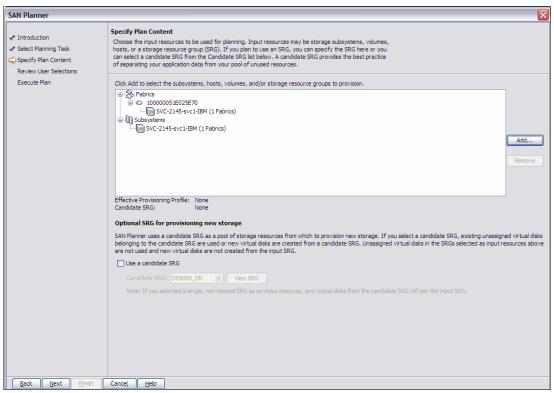


Figure 12-49 Specify Plan Content

5. For this example, we create a single VDisk with 5 GB of storage capacity and use a default **Space Only** workload profile (Figure 12-50). Click **Next**.

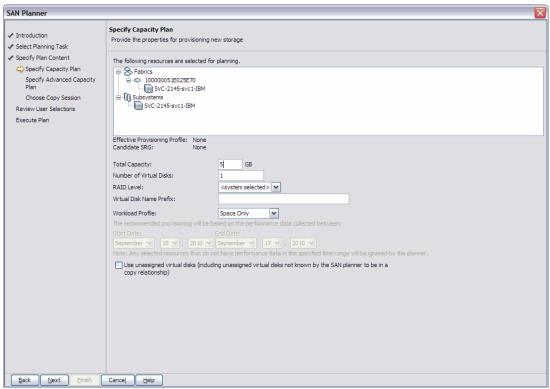


Figure 12-50 Specify Capacity Plan

6. Do not choose any advanced capacity plan options for Thin Provisioning or use SSDs. Leave these selections as the defaults (Figure 12-51) and click **Next**.

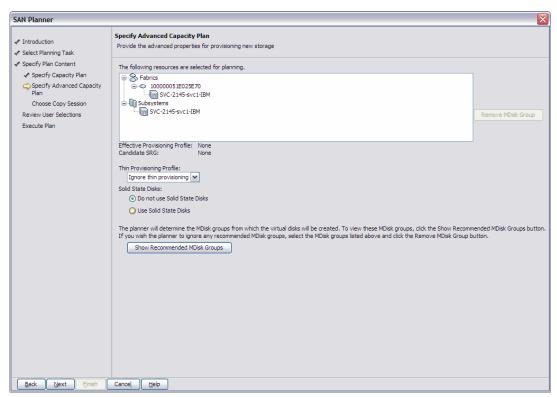


Figure 12-51 Specify Advanced Capacity Plan

7. Select **Create a new replication session** for the VDisk Mirroring relationship (Figure 12-52) and click **Next**.

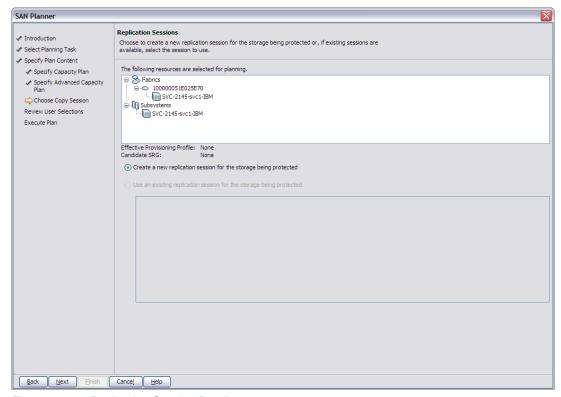


Figure 12-52 Replication Session Details

 On the session properties panel, select a session name and choose the Virtual Disk Mirroring option. Leave all other fields as the default options (Figure 12-53) and click Next.

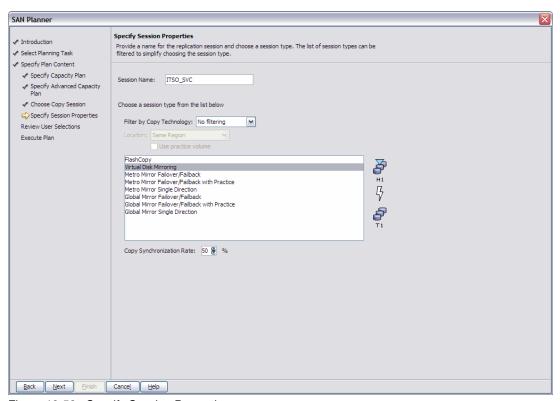


Figure 12-53 Specify Session Properties

9. Keep the secondary location default settings (Figure 12-54) and click Next.

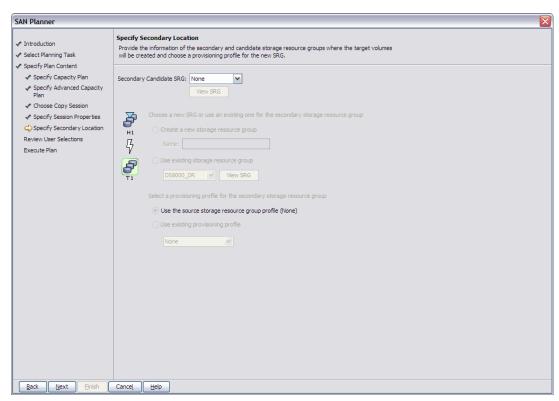


Figure 12-54 Specify Secondary Location

10. The SAN Planner provides a panel where you can review your selections. Ensure that all selections are accurate, and then click **Next** (Figure 12-55).

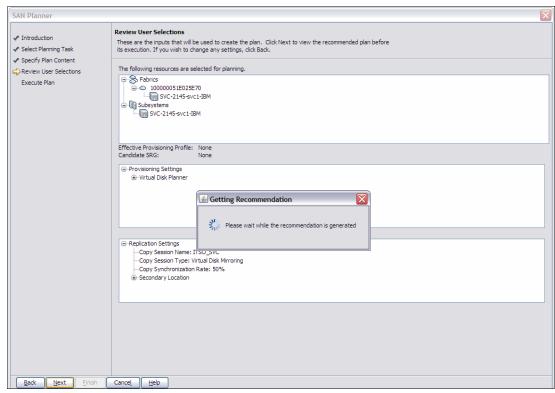


Figure 12-55 Review User Selections

The final panel lists the proposed changes (Figure 12-56). These include the primary Virtual Disk and the secondary Virtual Disk Mirror that will be created. Click **Finish** to save your plan and submit it immediately for execution using the option **Run Now**.

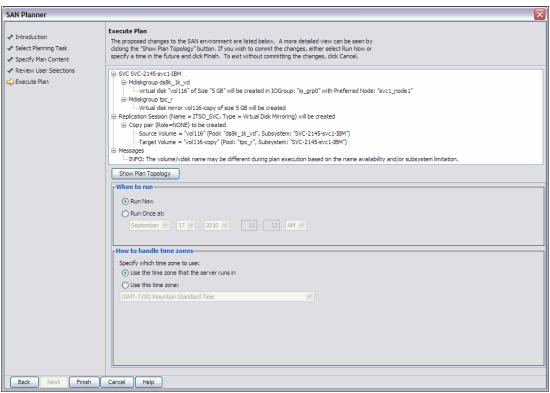


Figure 12-56 SAN Planner Confirmation Panel

11. The latest status is displayed in the job management panel (Figure 12-57).

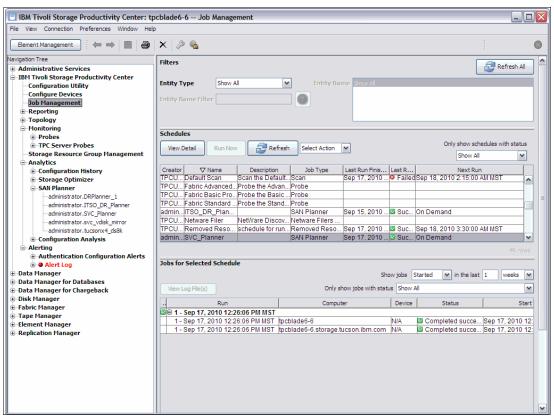


Figure 12-57 Job Management View for SAN Planner

After the job completes successfully, the virtual disk mirror details are in the SAN Volume Controller GUI (Figure 12-58).

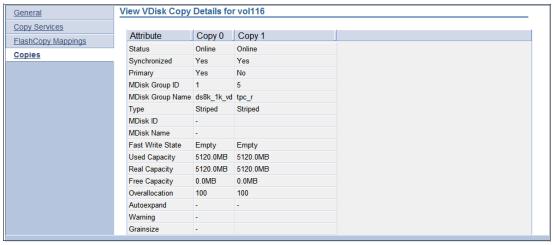


Figure 12-58 VDIsk Copy Details in the SAN Volume Controller GUI

12.9 SAN Planner with XIV replication

Tivoli Storage Productivity Center V4.2.2 introduces the SAN Planner support for the three types of XIV replication sessions:

- ► Snapshot
- Synchronous Metro Mirror failover/failback
- Asynchronous Global Mirror failover/failback

In this section we guide you through usage examples of those three replication plans for XIV.

12.9.1 Replication plan for XIV Snapshot

A Snapshot replication plan for XIV can be created together with a storage provisioning plan, including optional multipathing and zoning recommendations on the source resources (volumes or consistency groups), or in an isolated replication plan for existing storage resources.

In the following example we used the *Provision replication to existing storage* option, because provisioning was previously explained in other sections of this book. To create an XIV Snapshot plan, complete the following steps:

 To start the SAN Planner task, expand IBM Tivoli Storage Productivity Center → Analytics. Right-click SAN Planner, and select Create Plan (Figure 12-59).

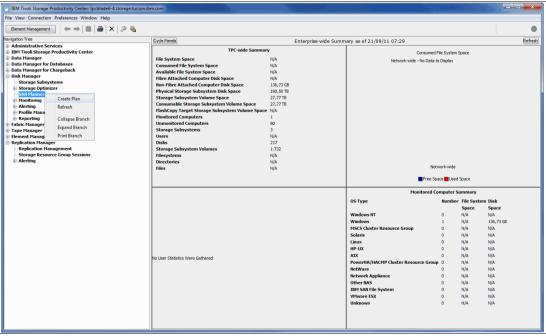


Figure 12-59 TPC Create Plan

2. In the Introduction page, click **Next** to open the Select Planning Task panel (Figure 12-60).

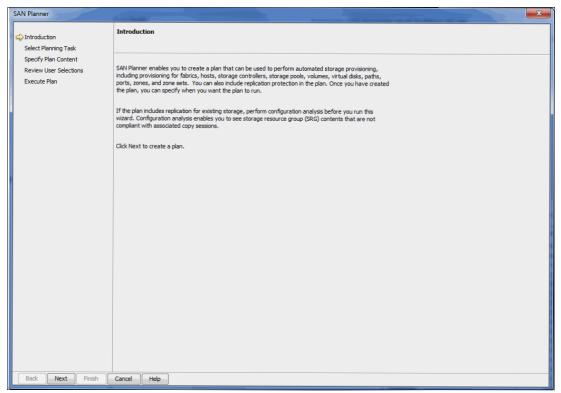


Figure 12-60 SAN Planner Introduction panel

3. Here, select **Provision replication to existing storage** to create a Snapshot from an existing volume (Figure 12-61). Another option in this panel allows you to provision new volumes and create a Snapshot from it in the same replication plan.

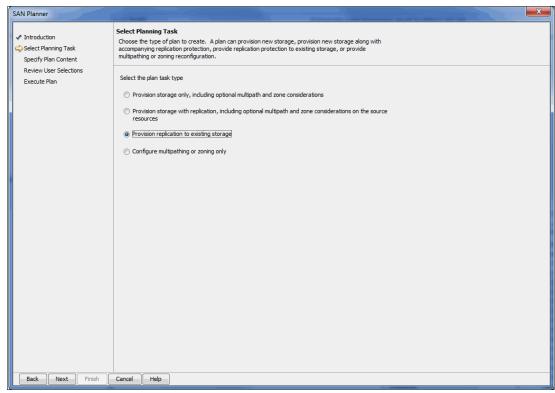


Figure 12-61 Select Planning Task panel - Provision replication to existing storage

4. On the Select Provision Plan Details panel, we choose to create a plan for existing Volumes, without multipath or zoning considerations (Figure 12-62). The other type, Virtual Disks, only applies to SVC or Storwize V7000.

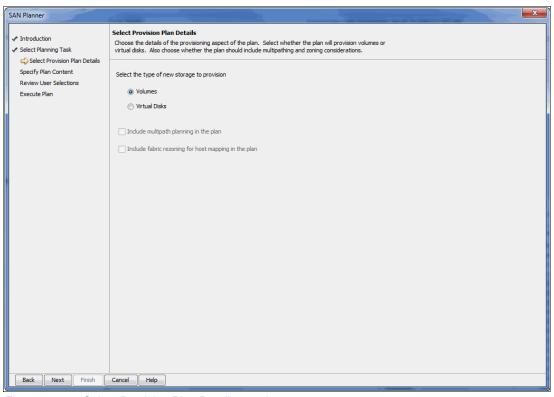


Figure 12-62 Select Provision Plan Details panel

5. For the plan content, we choose to manually select a single volume entity from the XIV to create our Snapshot plan (Figure 12-63). You also have the option to use a previously created Storage Resource Group containing multiple volumes as a Consistency Group to create a Snapshot replication plan on them, and select it with the *Use a candidate SRG* option in this panel instead. Click **Add** on the right side of the panel to go to the Planner Selection panel.

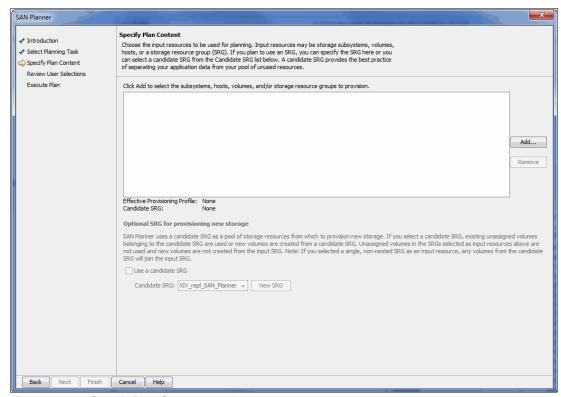


Figure 12-63 Specify Plan Content panel

6. The Planner Selection panel opens showing the *Available elements* where you can select the resources for your plan using the Topology Viewer. From the *Overview* section of the Topology Viewer, drill down to the target volume by double clicking the **Storage** box to see the current subsystems monitored by Tivoli Storage Productivity Center (Figure 12-64).

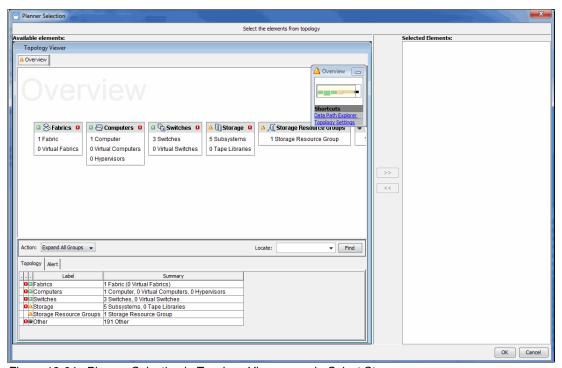


Figure 12-64 Planner Selection in Topology Viewer panel - Select Storage

7. In the L0:Storage section of the Topology Viewer, expand the **Subsystems** boxes to locate your target XIV by clicking the \oplus sign in the upper right corner of each of the Subsystems groups shown. When you find your target subsystem, double-click its icon to see the currently defined elements in that subsystem (Figure 12-65).

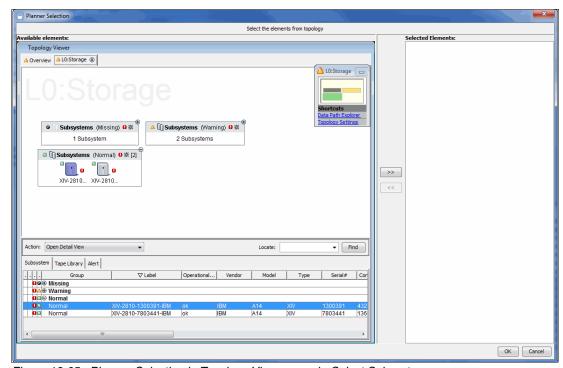


Figure 12-65 Planner Selection in Topology Viewer panel - Select Subsystem

Tip: If you know the subsystem's name, even partially, you can type it in the **Locate** field in the lower part of the Topology Viewer to find it more quickly. The resulting resource will be highlighted with a purple box around it in both the topology and tabular sections.

8. In the *L2:Subsystem* section of the Topology Viewer, expand the **Device** box by clicking the \oplus sign in its upper right corner to locate your target volume (Figure 12-66).

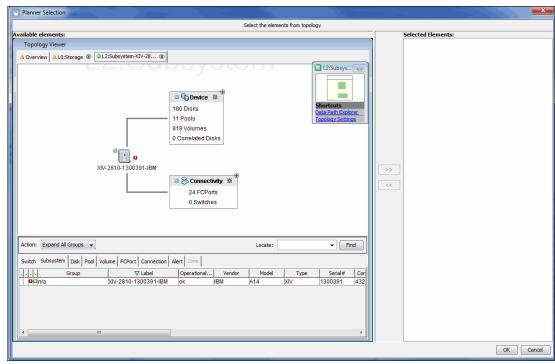


Figure 12-66 Planner Selection in Topology Viewer panel - Subsystem view

9. Locate the target volume in the various **Volumes** boxes that appear, which are grouped by Storage Pool within the selected XIV subsystem (Figure 12-67). Expand each box by clicking the \oplus sign in its upper right corner.

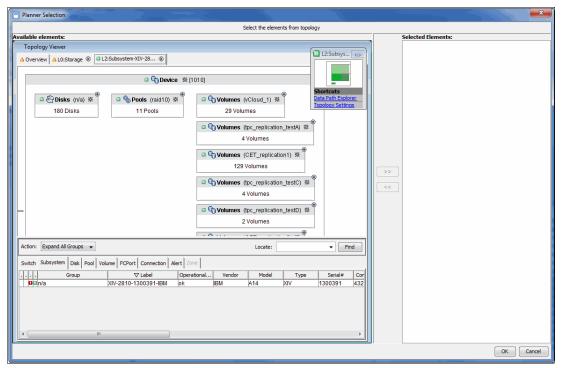


Figure 12-67 Planner Selection in Topology Viewer panel - Expand Devices

Tip: Again, if you know the volume's name, even partially, you can type it in the **Locate** field in the lower part of the Topology Viewer to find it more quickly. The resulting resource will be highlighted with a purple box around it in both the topology and tabular sections.

10. Select the target volume from which you want to create a Snapshot clicking its icon, then click the >> button in the middle of the panel to move the selected volume to the Selected Elements section to the right (Figure 12-68). Click OK to close the Planner Selection panel and go back to the Specify Plan Content panel.



Figure 12-68 Planner Selection in Topology Viewer panel - Select Volume

Tip: If you select more than one volume in this dialog, they will become part of a Consistency Group that contains one Copy Set for each source volume. When the Snapshot replication plan is executed, the Snapshot will be taken from the entire group of volumes in a consistent way.

11. When you have the plan content ready (Figure 12-69) click **Next** to go to the Replication Sessions panel.

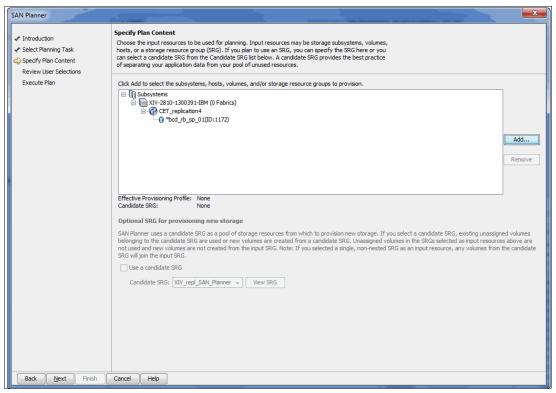


Figure 12-69 Specify Plan Content panel with selected volume

12.In the Replication Sessions panel you can create a new replication session or append the selected elements to an existing one. In our example, there are no previously created sessions, so we choose to create a new one (Figure 12-70). Click **Next** to go to the Specify Session Properties panel.

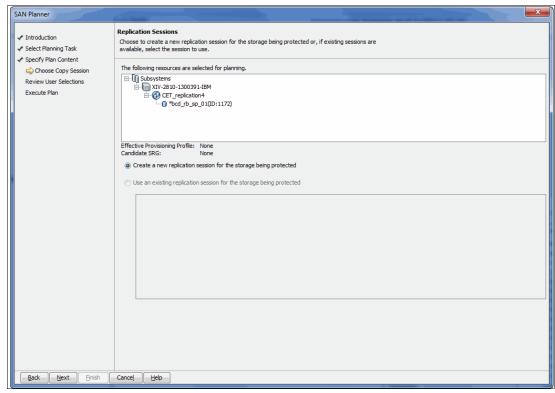


Figure 12-70 Replication Sessions panel

13.In the Specify Session Properties panel, we choose to create a Snapshot session. Notice that only the supported session types for the selected resources (XIV in our case) are listed (Figure 12-71). Click **Next** to go to the Review User Selections panel.

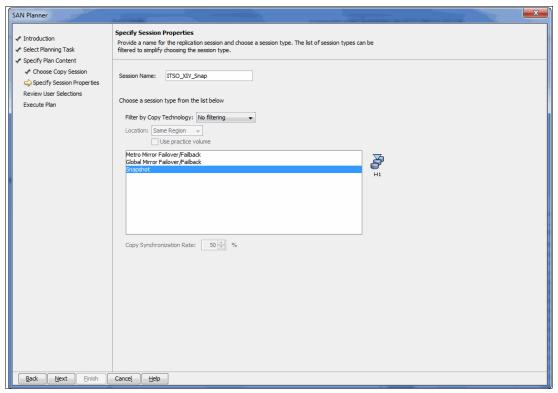


Figure 12-71 Specify Session Properties panel

14. In the Review User Selections panel, confirm that the desired elements are selected in the first field, and the session name and type are correct in the third field (Figure 12-72). Click Next to have the replication recommendations and all the required resources checked by the SAN Planner, and then go to the Execute Plan panel.

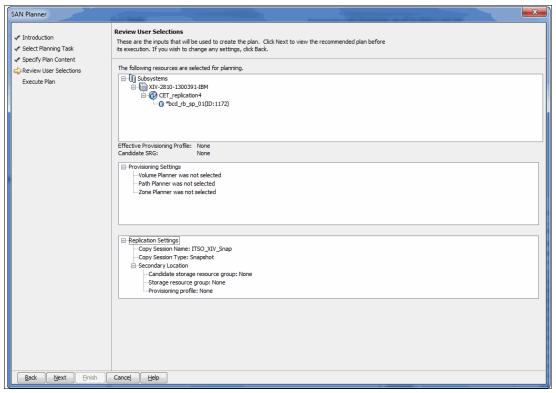


Figure 12-72 Review User Selections panel

Attention: If some conditions exist that might prevent the plan from being executed afterwards, you will receive alerts at this point, and the dialog cannot continue without those conditions being analyzed and fixed.

15. In the Execute Plan panel you can choose to run the plan immediately by selecting the **Run Now** option, or to schedule it to run in a more appropriate time using the **Run Once** at option and the applicable time zone handling (Figure 12-73). In our example, we choose Run Now. Click **Next** to enter a name for the plan to be saved.

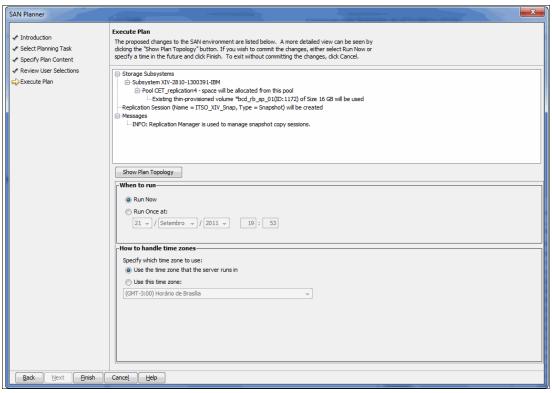


Figure 12-73 Execute Plan panel - Run Now

16. Type the name for your plan (Figure 12-74) and click **OK** to review the saved plan.

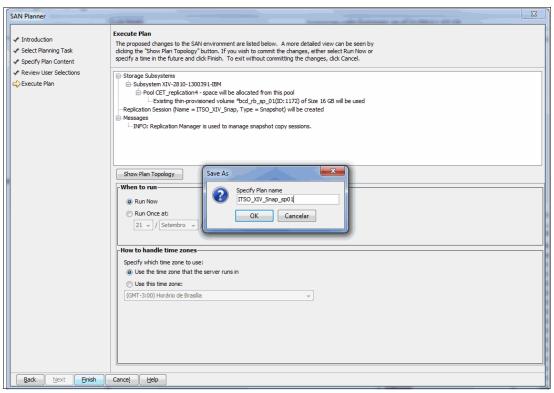


Figure 12-74 Specify Plan name

17.Select your saved plan under **Disk Manager** → **SAN Planner** in the Navigation Tree and click the **View Job History** button to review the results of your plan execution job (Figure 12-75).

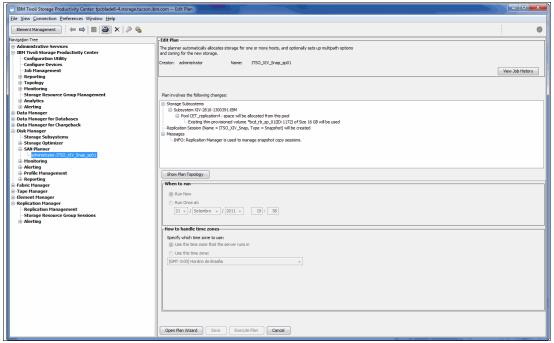


Figure 12-75 Review saved plan

18.Look for your job execution results to check if it was run successfully or not. You can use the column **Name** to find the job using the name of your plan. In our example, the job ended with no errors or warnings (Figure 12-76).

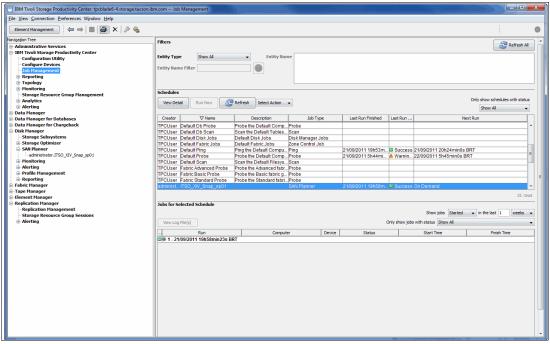


Figure 12-76 Review Plan job execution in Job Management panel

19. Because Snapshot sessions are not supposed to start without other conditions existing (for example, the affected volumes must be in a consistent state with the DBMS backup mode turned on), Snapshot replication plans executed by the SAN Planner only create the sessions in Tivoli Storage Productivity Center for Replication. They are not automatically started.

In order to have our plan implemented, we choose to go to the Tivoli Storage Productivity Center for Replication GUI to start our Snapshot session manually. Select **Replication Manager** \rightarrow **Replication Management** under the Navigation Tree and click the **Replication Sessions Overview** button (Figure 12-77).

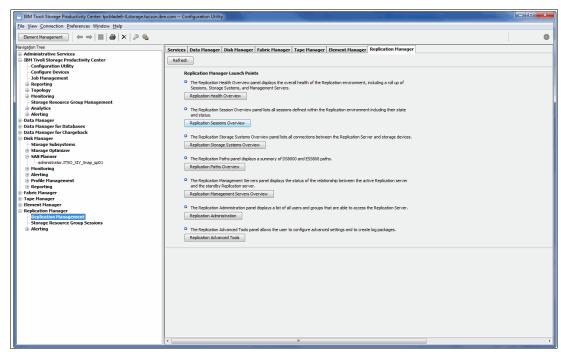


Figure 12-77 Launch Replication Manager to manually execute the Plan

20. The Tivoli Storage Productivity Center for Replication GUI opens in the Sessions panel. Select your Snapshot session from the list to take the appropriate action on it (Figure 12-78).

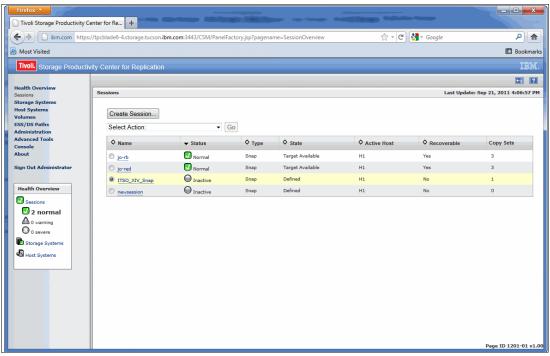


Figure 12-78 Select created session to be executed manually

21. From the **Select Action** pull down menu select **Create Snapshot** action to be executed (Figure 12-79).

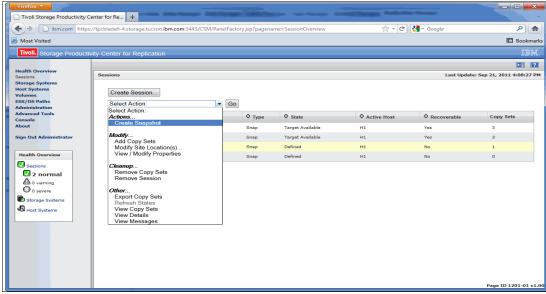


Figure 12-79 Create Snapshot in Tivoli Storage Productivity Center for Replication GUI

22.A confirmation message opens before the command is executed. Click **Yes** to have the Snapshot created (Figure 12-80).

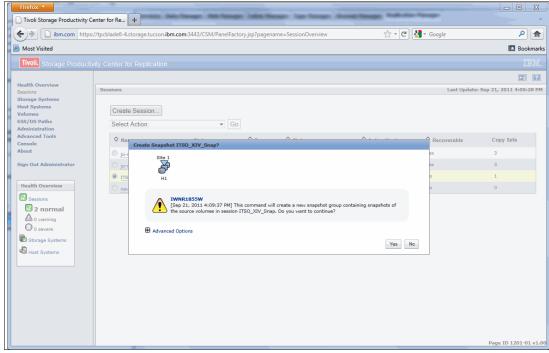


Figure 12-80 Create Snapshot confirmation message

23. Wait for the message that indicates your Snapshot was created successfully and the session status turns to **Normal** with a green check mark (Figure 12-81).

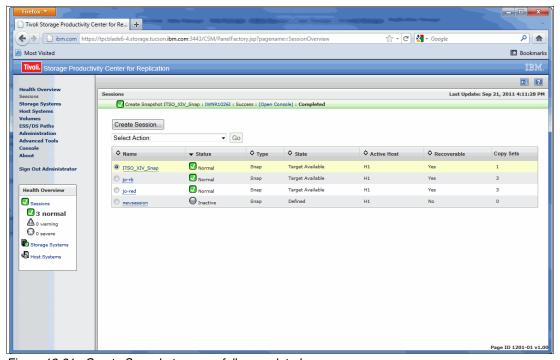


Figure 12-81 Create Snapshot successfully completed

24. As a last step in our example, we go to the XIV GUI to confirm that the Snapshot was really taken as expected in the *Volumes and Snapshots* panel (Figure 12-82).

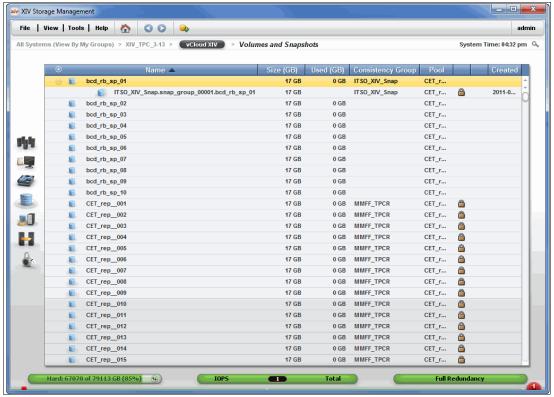


Figure 12-82 XIV GUI showing the Plan execution results

This completes our XIV Snapshot plan creation example using the SAN Planner.

12.9.2 Replication plan for XIV Metro and Global Mirror Failover/Failback

The other two types of replication sessions supported by the SAN Planner for IBM XIV in Tivoli Storage Productivity Center V4.2.2 are related to synchronous and asynchronous Remote Replication:

- Synchronous Metro Mirror Failover/Failback
- Asynchronous Global Mirror Failover/Failback

They can also be created together with a storage provisioning plan, including optional multipathing and zoning recommendations on the source resources (volumes or consistency groups), or in an isolated replication plan for already existing storage resources.

Similar to adding other subsystems in Tivoli Storage Productivity Center, for replication planning with XIV you must also add both subsystems to Tivoli Storage Productivity Center for Replication and ensure that connectivity is established between them prior to creating a plan.

In the following example we used the *Provision replication to existing storage* option, because capacity provisioning was previously explained in other sections.

To create an XIV Metro or Global Mirror Failover/Failback plan, complete the following steps:

 Add the storage devices to a candidate Storage Resource Group. Under IBM Tivoli Storage Productivity Center, select Storage Resource Group Management, and click Create (Figure 12-83).

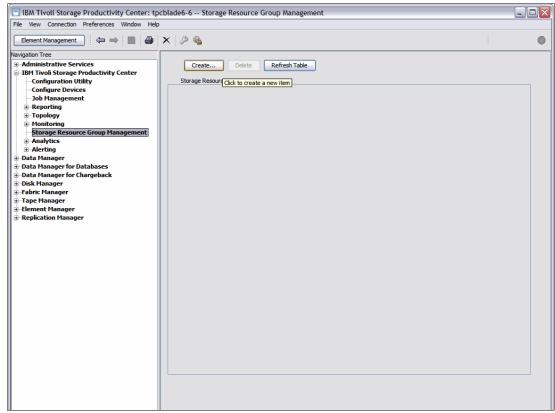


Figure 12-83 Storage Resource Group Management panel

2. The Create Storage Resource Group window opens (Figure 12-84). Type in a description for the SRG and click the **Add** button to the right to select elements to include in the SRG.

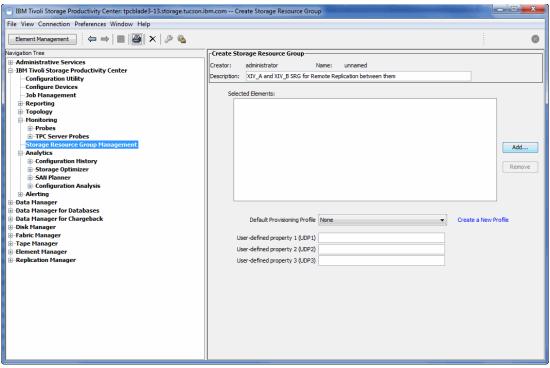


Figure 12-84 Create Storage Resource Group panel

3. In this example, we add two XIVs to the candidate group (Figure 12-85) moving them to the Selected Elements in the right side of the panel. Click **OK** to go back to the Create Storage Resource Group panel.

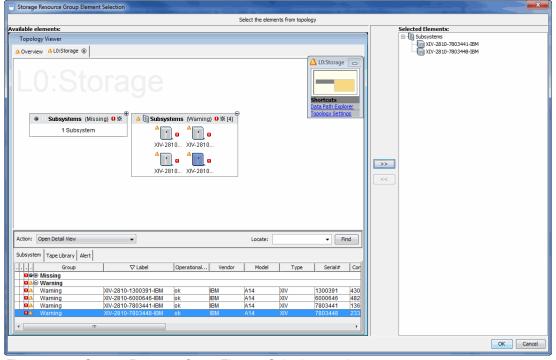


Figure 12-85 Storage Resource Group Element Selection panel

4. In the Create Storage Resource Group panel, save the Storage Resource Group by either clicking the disk icon or **File** → **Save** (Figure 12-86).

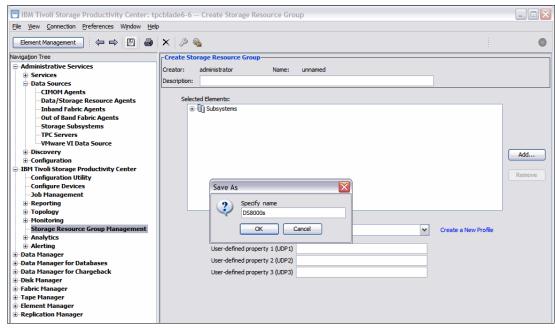


Figure 12-86 Specify SRG name to be saved

 After you have created the Storage Resource Group, you can start the SAN Planner task. Expand IBM Tivoli Storage Productivity Center → Analytics menu, right-click SAN Planner, and select Create Plan (Figure 12-87).

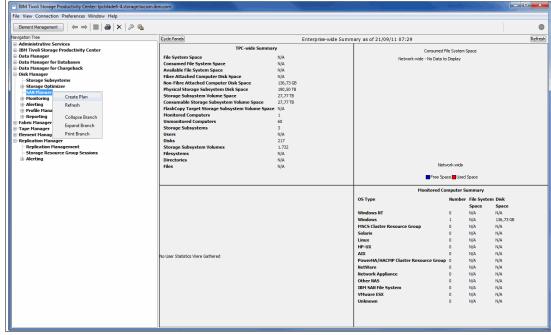


Figure 12-87 TPC Create Plan

6. In the Introduction page, click **Next** to open the Select Planning Task panel (Figure 12-88).

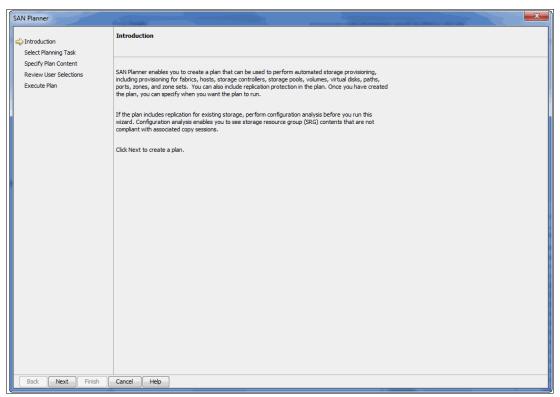


Figure 12-88 SAN Planner Introduction panel

7. Here, select Provision replication to existing storage to create a Remote Replication plan for an existing volume (Figure 12-89). You can also provision new volumes and create a Metro or Global Mirror from it in the same replication plan. In our example, we choose to create a replication plan for a set of three existing volumes in our Primary XIV. Click Next to go to the Specify Plan Content panel. See Figure 12-93 on page 522 for a list of the three volumes.

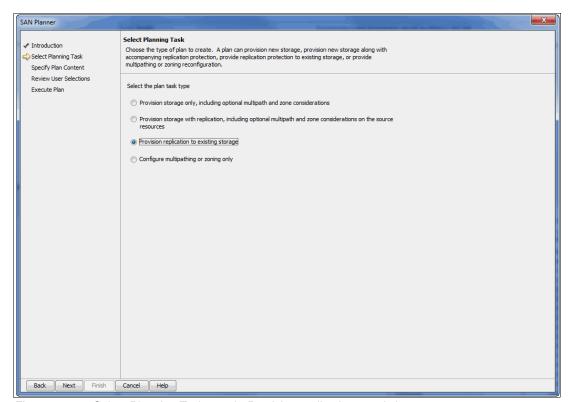


Figure 12-89 Select Planning Task panel - Provision replication to existing storage

8. On the Select Provision Plan Details panel, we choose to create a plan for volumes, without multipath or zoning considerations (Figure 12-90).

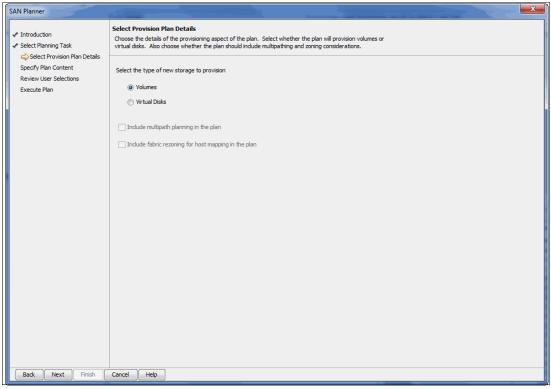


Figure 12-90 Select Provision Plan Details panel

9. For the plan content, we choose to select the Storage Resource Group we created before, containing the two XIV systems involved in the Remote Replication. Click **Add** in the right side of the panel to start the Planner Selection panel dialog (Figure 12-91).

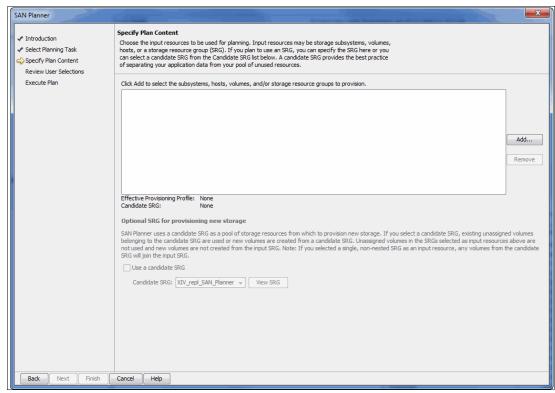


Figure 12-91 Specify Plan Content panel

10. The Planner Selection panel opens showing the *Available elements* from where you can select the resources for your plan using the Topology Viewer. From the *Overview* section of the Topology Viewer, drill down to the Storage Resource Group you want to include in the replication plan. Select the target SRG by clicking its icon. Click the >> button and move it to the Selected Elements to the right of the panel (Figure 12-92).

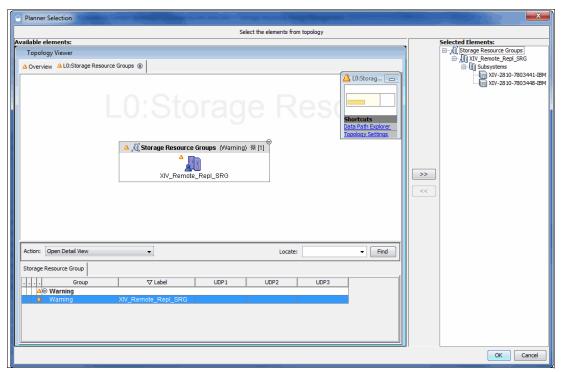


Figure 12-92 Specify Plan Content - select SRG

11. Because our plan is going to provision replication for existing volumes, we need to include them in the Selected Elements as well. Navigate the Topology Viewer to the target volumes, select them and click the >> button to move them to the Selected Elements to the right (Figure 12-93). In our example, we choose to select three volumes to form a Consistency Group. Click **OK** to close the Planner Selection panel and go back to the Specify Plan Content panel.

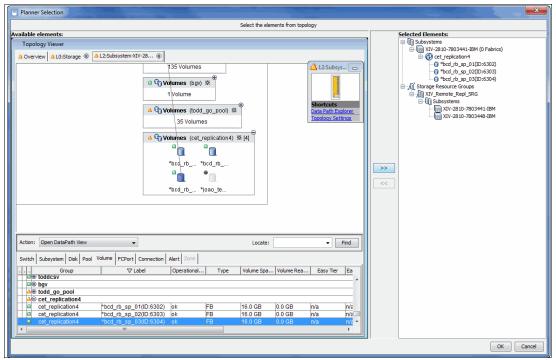


Figure 12-93 Specify Plan Content - select existing volumes to be replicated

Tip: If you know the name of the volumes, even partially, you can type it in the **Locate** field in the lower part of the Topology Viewer to find them more quickly. The resulting resource will be highlighted with a purple box around it in both the topology and tabular sections.

12. When you have the plan content ready (Figure 12-94) click **Next** to go to the Replication Sessions panel.

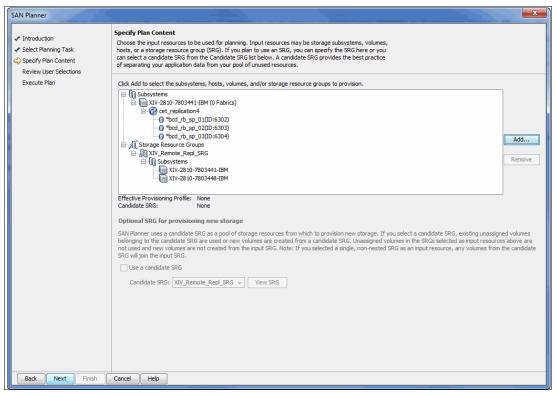


Figure 12-94 Specify Plan Content panel with selected volumes

13.In the Replication Sessions panel you can create a new replication session or append the selected elements to an existing one. In our example, there are no previously created sessions, so we choose to create a new one (Figure 12-95). Click **Next** to go to the Specify Session Properties panel.

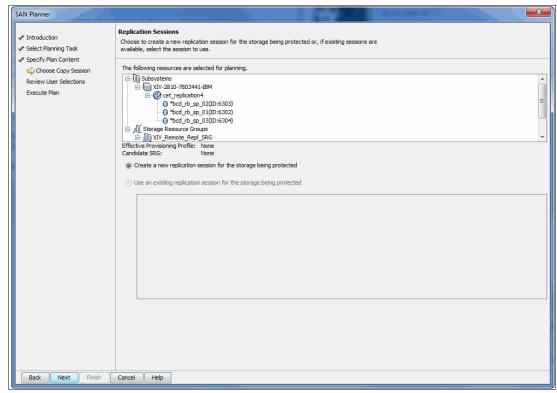


Figure 12-95 Replication Sessions panel

14. In the Specify Session Properties panel we choose to create a Metro Mirror Failover/Failback session. Notice that only the supported session types for the selected resources (XIV in our case) are listed (Figure 12-96). Click **Next** to go to the Specify Secondary Location panel.

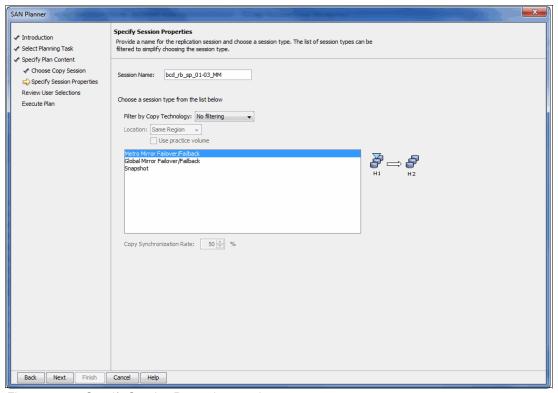


Figure 12-96 Specify Session Properties panel

Replication: The other type of remote replication session available for XIV is *Global Mirror Failover/Failback*, which is an asynchronous Remote Replication function, whereas Metro Mirror is synchronous. Although these two replication scenarios are different, with particular considerations for each one, the process of creating and executing a replication plan with SAN Planner is the same for both. Select the appropriate option that applies to your replication environment.

15. In the Specify Secondary Location panel, we choose to create a new target SRG to accommodate the Metro Mirror Failover/Failback session (Figure 12-97). Type in the chosen target SRG name to be created and click **Next** to go to the Review User Selections panel.

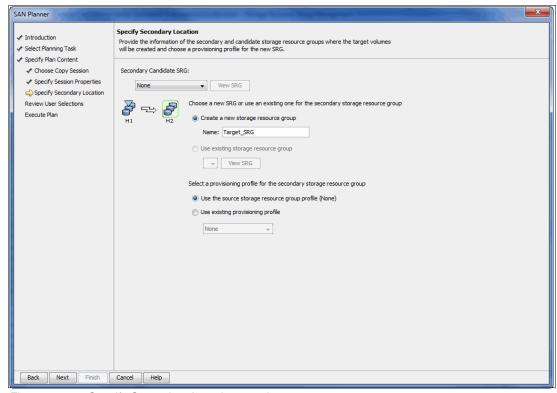


Figure 12-97 Specify Secondary Location panel

16. In the Review User Selections panel confirm that the desired elements are selected in the first field, and the session name and type are correct in the third field (Figure 12-98). Click Next to have the replication recommendations. All the required resources are checked by the SAN Planner.

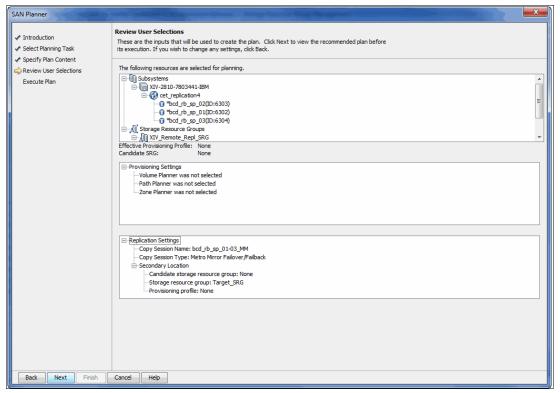


Figure 12-98 Review User Selections panel

Attention: If some conditions exist that might prevent the plan from being executed afterwards, you will receive alerts at this point, and the dialog cannot continue without those conditions being analyzed and fixed.

17. In the Execute Plan panel, you can choose to run the plan immediately by selecting the **Run Now** option, or to schedule it to run in a more appropriate time using the **Run Once** at option with the applicable time zone handling (Figure 12-99). In our example, we choose **Run Now**. Click **Next** to enter a name for the plan to be saved.

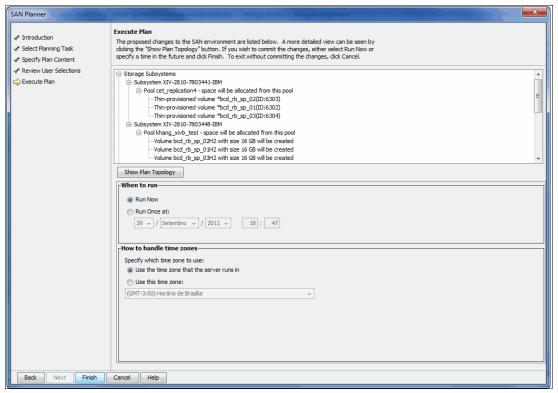


Figure 12-99 Execute Plan panel - Run Now

18. Type the name for your plan (Figure 12-74) and click **OK** to review the saved plan execution status.

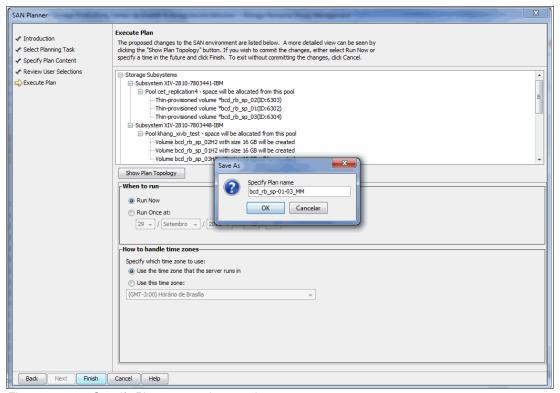


Figure 12-100 Specify Plan name to be saved

19. Using the **Job Management** option in the Navigation Tree, look for your job execution results to check if the job was run successfully or not. You can find it by the name of your plan in column Name or using other filters available in the Job Management panel. In our example, the job ended with no errors or warnings (Figure 12-101).

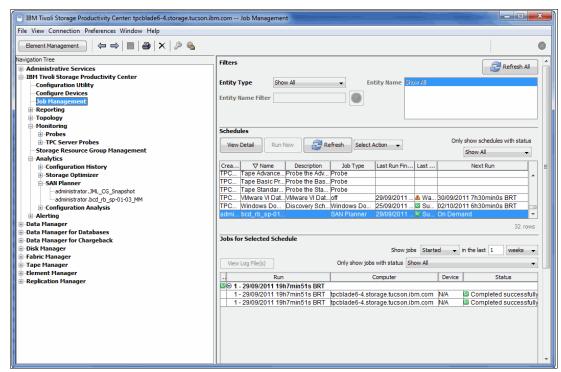


Figure 12-101 SAN Planner job execution completion

20. Remote Replication sessions have requirements to be met before they can be started (for example, the replication links must be established and up and the two involved subsystems are paired). Remote Replication plans executed by the SAN Planner only create the sessions in the Tivoli Storage Productivity Center for Replication, and they go into a *Prepared* status. They are not automatically started.

In order to have our plan implemented, we choose to go to the Tivoli Storage Productivity Center for Replication GUI to start our Metro Mirror session manually. Select **Replication Manager** \rightarrow **Replication Management** under the Navigation Tree and click the **Replication Sessions Overview** button (Figure 12-102).

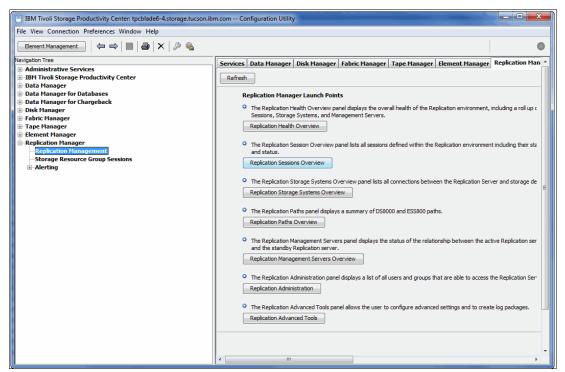


Figure 12-102 Replication Management - Replication Sessions Overview

21. The Tivoli Storage Productivity Center for Replication GUI opens in the Sessions panel (Figure 12-103). Click the name of the session to see its details.

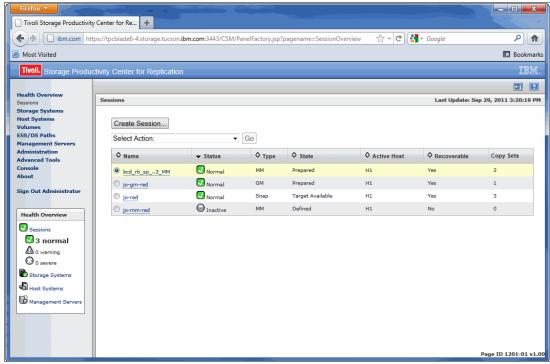


Figure 12-103 Tivoli Storage Productivity Center for Replication Sessions panel

22. Review the Remote Replication session details (Figure 12-104).

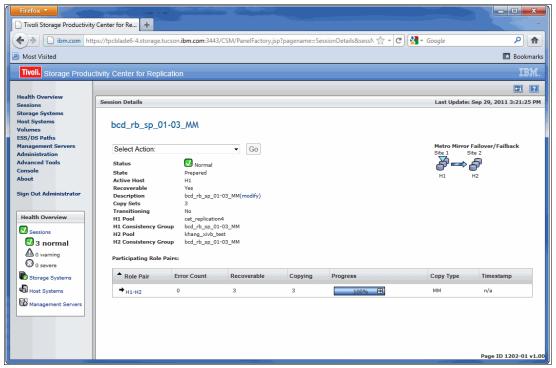


Figure 12-104 TPC for Replication Session Details panel

23. From the **Select Action** pull-down menu, select the **Start H1->H2** action to be executed (Figure 12-105).

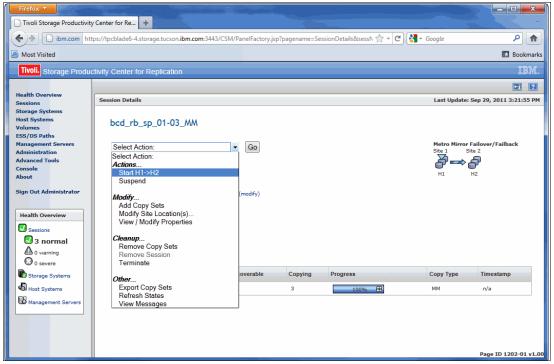


Figure 12-105 TPC for Replication Start H1-H2 Action menu

24.A confirmation message opens before the command is executed. Click **Yes** to have the Remote Replication from Primary to Secondary started (Figure 12-106).

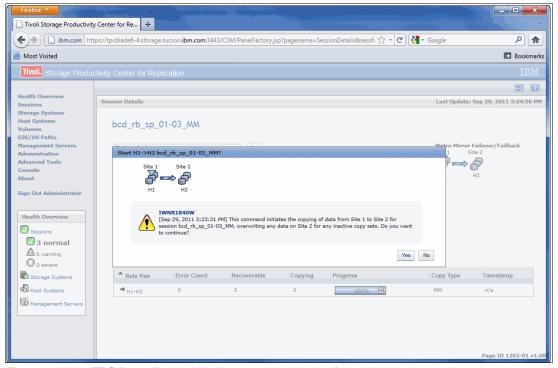


Figure 12-106 TPC for replication, H1-H2 remote replication from primary to secondary start

25. Wait for the message indicating that your Remote Replication started successfully and the session status turns to **Normal** with a green check mark. Go to the Primary XIV GUI to confirm that the Remote Replication was really started and are in a *Synchronized* status as expected in the Remote Mirroring panel (Figure 12-107).



Figure 12-107 Primary XIV GUI Remote Mirroring panel

26.Go to the Secondary XIV GUI to confirm that the target volumes in the Remote Replication session are in a *Consistent* status under the same Consistency Group as expected in the Remote Mirroring panel (Figure 12-108).

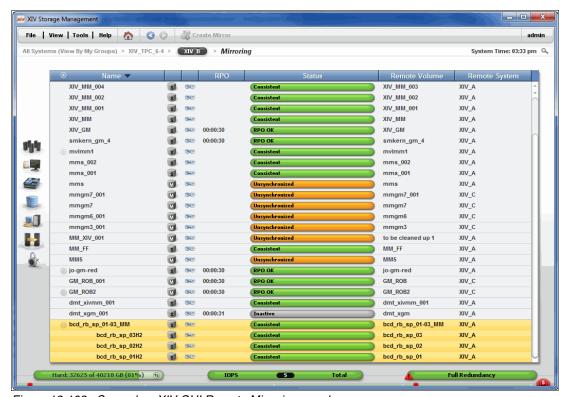


Figure 12-108 Secondary XIV GUI Remote Mirroring panel

This completes our XIV Remote Replication plan creation example using the SAN Planner.



Job Management panel

The Job Management panel is new in Tivoli Storage Productivity Center V4.2 after customer feedback. The panel is the central place to manage any job within Tivoli Storage Productivity Center. In previous versions, you had to look at the Navigation Tree to find information about completed or running jobs. That cumbersome task is now consolidated in the Job Management panel.

This chapter explains the Job Management panel and how it is invoked, and includes steps of a common task to help you become familiar with using this feature.

Additionally, this chapter describes the most important terms that are used for managing jobs in Tivoli Storage Productivity Center.

Terminology used: Certain system menus in this chapter issue "recommendations" that are meant to serve as guidelines for your choices.

13.1 Background

Previously, the Tivoli Storage Productivity Center GUI had limitations and issues in the way schedules, runs, and jobs were located and displayed in the GUI:

- ► Given a particular device, there was not an easy way to identify all schedules in which it participated.
- ► There was not an easy way to determine what was happening in the Tivoli Storage Productivity Center server at any given moment.
- ► There was not an easy way to determine what jobs within Tivoli Storage Productivity Center were having problems, and therefore what portion of the environment might not have been monitored to its fullest potential.
- ▶ Any control activities invoked through APIs (Tivoli Provisioning Manager Workflow integration) or CLIs were not visible in the Tivoli Storage Productivity Center GUI.
- ► Any mini-probes or activities initiated by some event in the storage network were not visible in the Tivoli Storage Productivity Center GUI

The new Job Management panel addresses these issues by consolidating the management job schedules and jobs into a central panel.

Filter help to reduce the information displayed so that you can focus on a single or a few devices of interest. For example, you can use Storage Resource Groups to view only schedules with the devices associated with a particular group.

Recommendations will be generated when Tivoli Storage Productivity Center is not fully used to monitor the storage environment. Typically these are to add a performance monitor job for a device.

13.2 Job Management terms

This section describes three terms and how they relate to each other:

- ► Schedule
- ► Run
- ▶ Job

These terms can sometimes be used in a confusing way. For example, context menu commands exist for schedules in the Navigation Tree to Refresh Job List and Update Job Status, as seen in Figure 13-2 on page 540.

13.2.1 Schedule

A *schedule* is what we commonly refer to as a job "definition." The Default Probe created during the installation of Tivoli Storage Productivity Center is an example of a schedule.

13.2.2 Run

A *run* (or job run) is a single invocation of an entire schedule. One schedule might (and almost always does) have multiple runs. Each run has an associate number (the run number) that begins with 1, and increments each time the schedule is invoked again. In versions prior to V4.2, the Tivoli Storage Productivity Center GUI typically showed the run number followed by the date and time of that run. For an example, see Figure 13-3.

13.2.3 Job

A *job* is a unit of work within a schedule run; the degree of complexity for this unit of work varies by implementation. For example, with a subsystem probe schedule, one job is created per subsystem in the schedule. One run of a schedule might have multiple jobs.

For certain types of schedules, each run includes two wrapping jobs, which usually log the start and end of the other job or jobs that are executed during this run.

13.3 Job Management changes

With Tivoli Storage Productivity Center V4.2, certain changes have been made to the way a user manages jobs. In previous versions, the user would select the Navigation Tree item to do the following functions:

- ► Create a new schedule.
- ► View details and modify existing schedule.
- ▶ Delete a schedule.
- ▶ View the status of the last schedule.
- ► Stop the current run of a schedule (not available for all jobs).
- ► Open logs of the jobs of a run.
- View the schedule/job history.

Most of these functions are also available through the new Job Management panel, except for creating new schedules.

Displaying job runs as nodes in the Navigation Tree has been removed from Tivoli Storage Productivity Center. Figure 13-1 compares how the Navigation Tree appears in separate Tivoli Storage Productivity Center versions. Additionally highlighted in this figure is the location of the new Job Management panel in the Navigation Tree.

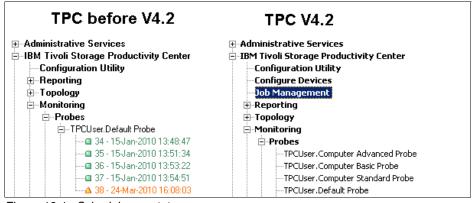


Figure 13-1 Schedule run status

The context menu entries have been adjusted, as shown in Figure 13-2.

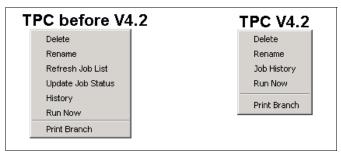


Figure 13-2 Schedule Context menu

If you now click **Job History** (previously simply called History), the Job Management panel opens.

13.3.1 Default jobs

Prior to Tivoli Storage Productivity Center V4.2, Navigation Tree entries existed for tasks that a user invoked but that were run in the background, for example provisioning of a LUN. In this example, there is a new status entry for the LUN provision job (Job number 5), under Disk Manager \rightarrow Monitoring \rightarrow Jobs, as shown in Figure 13-3.

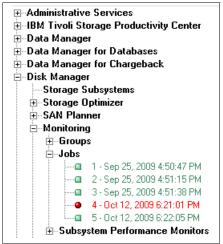


Figure 13-3 Background jobs before V4.2

Starting with Tivoli Storage Productivity Center V4.2, the *Jobs* Navigation Tree item has been retained just as the other schedule entries. However, the only action that you can do is to open the context menu and go to the Job Management panel by clicking the *Jobs* menu item, as seen in Figure 13-4.

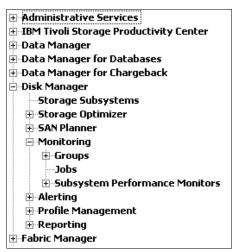


Figure 13-4 Background jobs starting with V4.2

Depending on the Disk manager or Fabric manager, one of the following schedules is preselected:

- ► Default Disk Jobs
- Default Fabric Jobs

13.3.2 CLI and event-driven jobs

For any API, CLI, and background driven activity, a new pseudo schedule named *CLI and Event Driven Jobs* has been added to the list of schedules within Tivoli Storage Productivity Center. One typical example of such a job is the *probe* that will be started when that invocation of the Configure Devices wizard completes.

The user can also see more information about activities initiated by external programs, for example, Tivoli Provisioning Manager Workflows.

13.4 Job Management panel explained

This section shows how the new panel looks and describes functions available in the panel. You can open the Job Management panel in any of the following ways:

▶ Icon bar:

A new icon is available on the icon bar of the Tivoli Storage Productivity Center GUI window. The new icon (highlighted in Figure 13-5) can be used to access the Job Management panel.



Figure 13-5 Job Management Icon

Navigation Tree:

You can select **Job Management** from the Navigation Tree, as shown on the right side of Figure 13-1 on page 539.

Context menu:

You can click **Job History** on the context menu to open the Job Management panel, with the schedule that the context menu was opened for being preselected. See the right side of Figure 13-2 on page 540.

Number of panels open at a time:

- Although you can open the Job Management panel from multiple places within the Tivoli Storage Productivity Center GUI, only one panel can be open at any time. Opening the panel a second time resets the already opened panel, and you lose your current selections.
- ► The Log File Viewer panels can each display multiple log files in tabs, and you can have multiple Log File Viewer panels open at the same time. Figure 13-12 on page 551 shows the Log File Viewer panel.

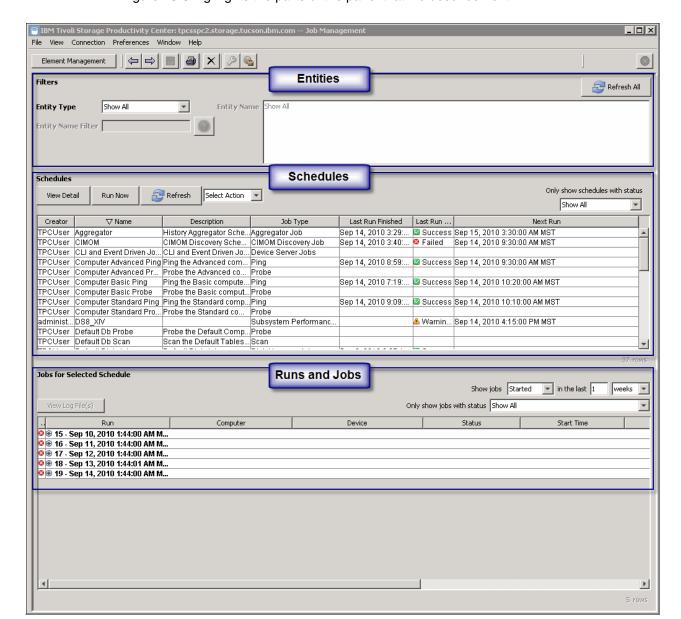


Figure 13-6 highlights the parts of the panel that we describe next.

Figure 13-6 Panel sections

The highlighted areas in Figure 13-7 indicate the parts of the Job Management panel that you can use to filter the amount of information shown. Both the Schedules section and the Runs and Jobs section display rows in each of their tables (shown only for the Runs and Jobs section).

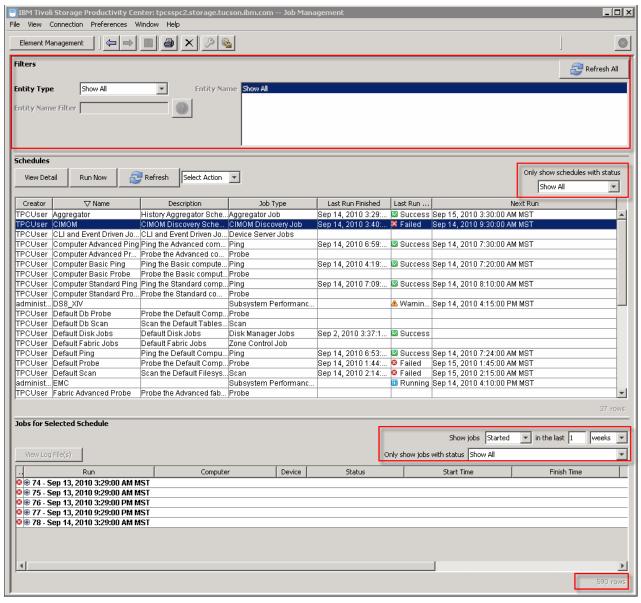


Figure 13-7 Filters

Usually you start from the top of the panel in the Entities section by selecting or filtering what you want to be displayed in the middle Schedules section. The middle section does not offer so many filters, but most of the time you reduce the amount of information displayed here by selecting only a few devices in the Entity Name field in the Entities section. Selecting a schedule will show the Runs and Jobs in the third and bottom section of the panel from which you can further use filters, and finally open the log files.

Although this approach might seem complex, you soon realize that this panel consolidates information about jobs for a single device or groups of devices, including discoveries, probes, performance monitor jobs, and more.

13.4.1 Entities details

In this section of the panel, you can filter on what you want to see in the other two sections (Schedules and Runs and Jobs), which are displayed below the Entities section. Although you will mostly select devices here, Monitoring Groups and Storage Resource Groups can also be selected. Therefore, this section was given the generic term of *Entities section*.

First, select one of the following items in the Entity Type field:

- ► Show All
- Storage Subsystem
- ▶ Fabric
- ► Computer
- Switch
- ► Tape Library
- Monitoring Group
- Storage Resource Group

After you have selected the type, you can either use the Entity Name Filter on the left side or form the list on the right side, to further reduce what is displayed in the Schedules section.

Tip: If you select a device in this part of the panel, you might see a recommendation, as shown in Figure 13-13 on page 552, which tells you that you can and should further set up that device. Usually these recommendations are for setting up performance monitors, which are not part of the Configure Devices Wizard setup, but you must determine in advance whether you have the correct license that allows you to create and run Performance monitors.

13.4.2 Schedules details

The Schedules section displays all schedules that are related to the devices that have been selected in the Entities section. From here, you can invoke functions by highlighting a schedule and either clicking one of the tabs (View Detail, Run Now, or Refresh) or you can also select an action from the Select Action drop-down menu. The actions list depends on the schedule that has been highlighted, so for example the default schedules do not have a rename or delete action associated with them.

You can either use the Only show schedules with status menu in the upper right of the Schedules section or sort the schedules by clicking the heading of one of the columns.

Tip: The labels of the Job Type Discovery jobs are not always as descriptive as they could be, so that sorting on the Job Type column does not group all Discovery jobs nicely together.

Prior to Tivoli Storage Productivity Center V4.2, the Description fields were never displayed on any panel other than the Schedules panel. With V4.2, you can use the Description field to work around the Discovery Job type limitation. You can use the Description field of those jobs and edit the Description to always start with Discovery so that you can sort on this column instead of the Job Type.

13.4.3 Runs and Jobs details

After you have selected a schedule, the runs of this schedule are displayed in the last section located at the bottom of the Job Management panel. Here, more filters are available to reduce the number of runs and jobs to be displayed. At this point, the filters are more oriented at the time that the runs and jobs were executed or started; the filtering for other fields such as computer or device was already done in the Entities section of the panel.

When you open the details of a run by clicking the plus sign (\oplus) located in the left side of the run name, you see the jobs of that run, as seen in Figure 13-11 on page 550 for run number 8. You can also open multiple runs, and select multiple jobs so that clicking **View Log File(s)** will open one tab for each of the selected jobs with the corresponding log file content as shown in Figure 13-12 on page 551.

13.5 Example

For this example, we have created a Storage Resource Group (SRG) containing one SVC and the back-end device that this SVC had attached at this time.

13.5.1 List Schedule log files

The example steps are as follows:

1. When you open the Job Management panel, no filter has been specified in the Entity Type field of the Filters section, so all schedules are shown. Select the Storage Resource Group as the Entity Type. See Figure 13-8.

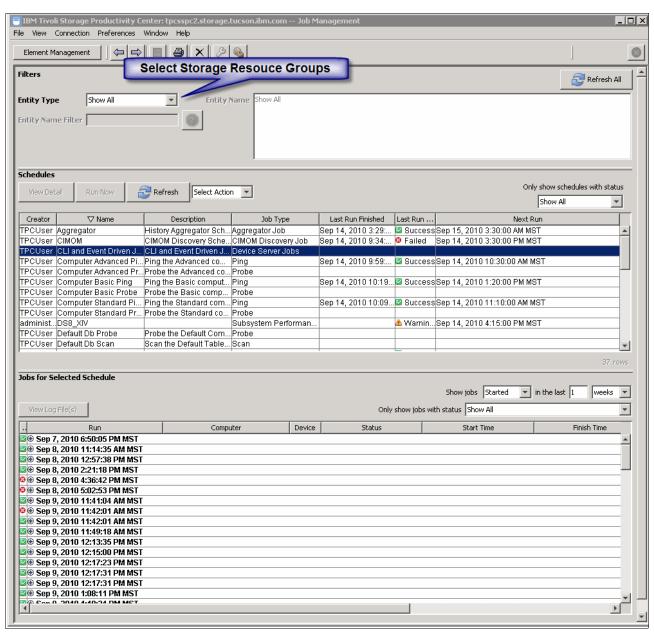


Figure 13-8 Select Storage Resource Group

2. After you have selected the entity type (here Storage Resource Group), only the defined SRGs are displayed (Figure 13-9).

We enter a name filter to reduce the number of SRGs displayed. We entered a name in the Entity Name Filter field to further reduce the number of Storage Resource Groups displayed. Notice that with each letter entered in the Entity Name Filter, the list on the Entity Name field is refined.

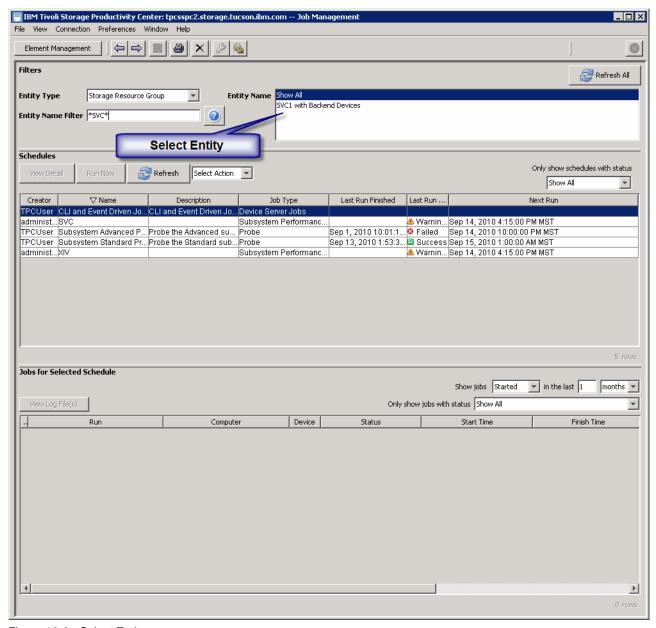


Figure 13-9 Select Entity

3. In the Entity Name list on the right, you can change the selection from Show All, and select one or more of the listed entities. In our example, this had no effect, because the list only included one entity, therefore Figure 13-10 does not look different.

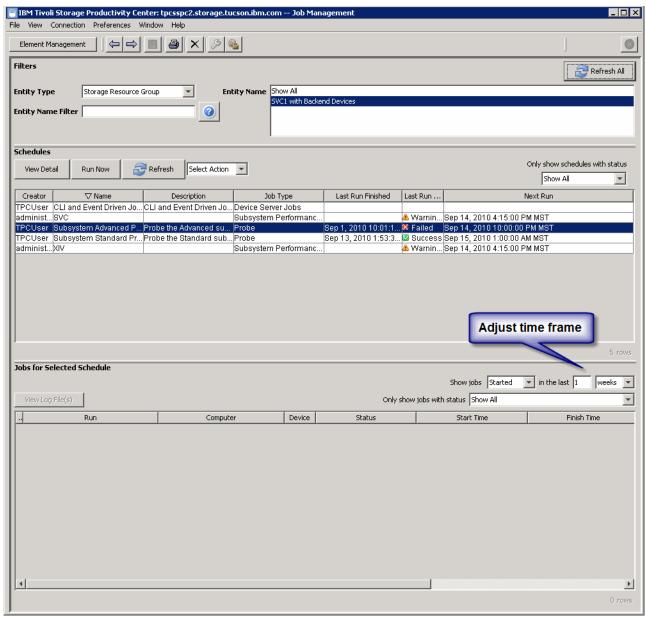


Figure 13-10 Select a schedule and adjust time frame

4. We focus our attention on the failed scheduled Probe job and selected that entry. In the bottom section of the panel nothing changed, because the last run was more than a week ago, so we changed the Show Jobs filter to show entries for the last month (Figure 13-11).

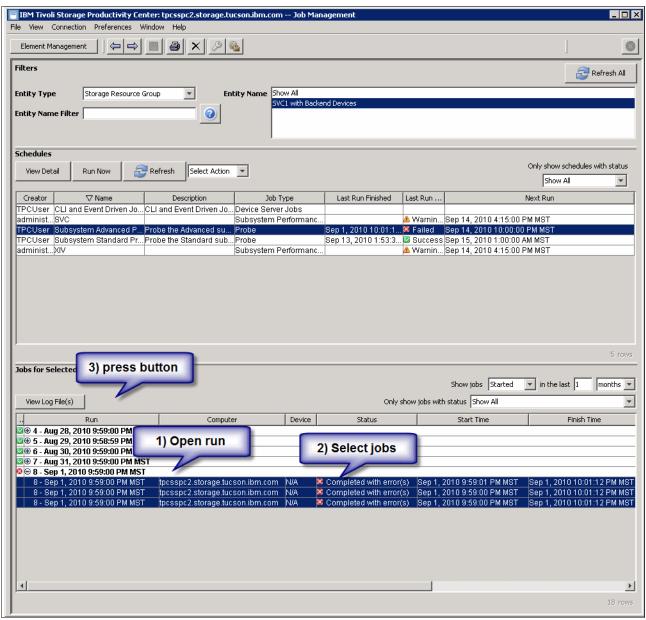


Figure 13-11 Select log files and open Log File Viewer

- 5. We use the following steps:
 - a. Open the failed run from the Run list
 - b. Select the three jobs that belong to that run
 - c. Click View Log File(s) at the top of the Run list.

Figure 13-12 shows the Log File Viewer with three tabs, one for each of the selected Job log files from Figure 13-11 on page 550.

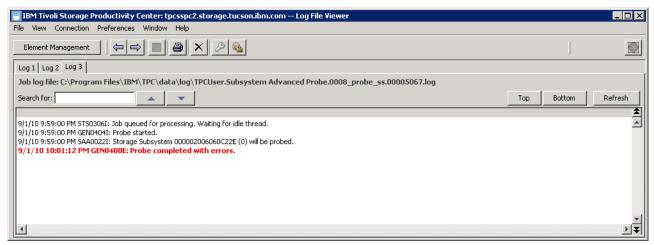


Figure 13-12 Log File Viewer panel

13.5.2 View and act on Tivoli Storage Productivity Center recommendation

After we added the SVC from the previous example to Tivoli Storage Productivity Center, we did not have a performance monitor or Storage Resource Group setup, so when we selected the SVC by using the Storage Subsystem Entity filter, Tivoli Storage Productivity Center gave the recommendation shown in Figure 13-13.

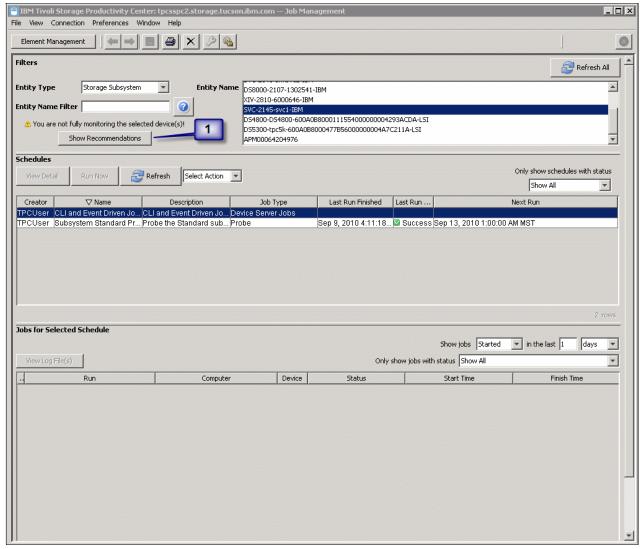


Figure 13-13 Recommendation displayed

When you click **Show Recommendations** (Figure 13-13), details about the recommendation are displayed (Figure 13-14).

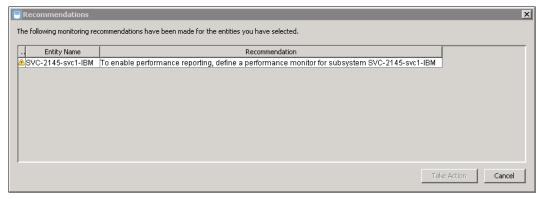


Figure 13-14 Recommendation details

In our example, only one device is selected so only one recommendation is displayed. Figure 13-15 is taken at a different time in the same environment, so there are many more recommendations for various device types.

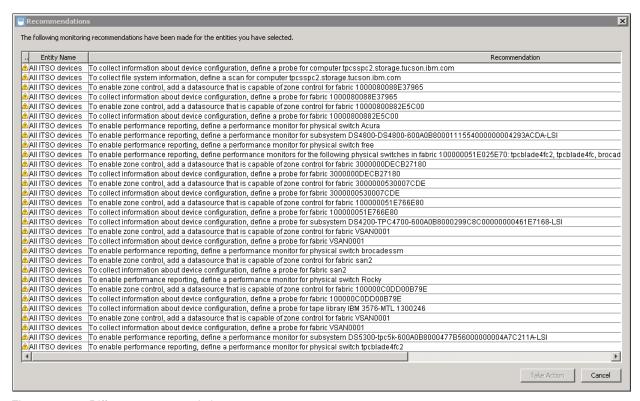


Figure 13-15 Different recommendations

If you want to implement recommendations, highlight them and click Take Action.

In this case Tivoli Storage Productivity Center opens the panel for creating a new performance monitor, but does not actually add the device to the monitor yet, as shown in Figure 13-16. So you still need to complete the action in that panel in order to create and save the required performance monitor.

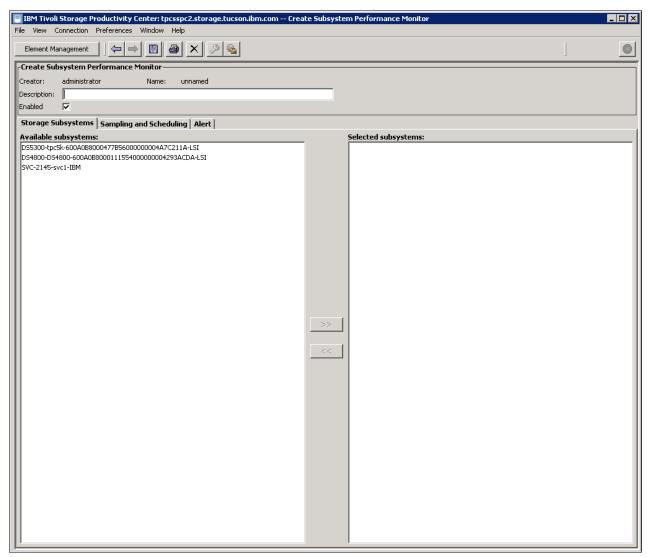


Figure 13-16 Take Action button opens the Performance Monitor Job panel



14

Fabric enhancements

IBM Tivoli Storage Productivity Center V4.2 fabric management capabilities improve on previous versions by adding support for new host bus adapters (HBAs), Converged Network Adapters (CNAs), and switch models, and enabling integration with Brocade DCFM and Brocade Network Advisor software. Also, Tivoli Storage Productivity Center V4.2 now offers basic Fibre Channel over Ethernet (FCoE) support.

In this chapter, we briefly describe these new features.

14.1 Fibre Channel over Ethernet (FCoE) support

The following capabilities are supported:

- ► Tivoli Storage Productivity Center V4.2 now includes toleration support of FCoE switches and CNA cards.
- ► Testing has been performed with Brocade 8000 and Cisco Nexus 5000 switches, and QLogic, Brocade, and Emulex CNAs.
- ► The Converged Enhanced Ethernet (CEE) or FCoE connectivity functions are not supported.
- ► Performance monitoring is supported on the Fibre Channel (FC) ports, but not on the FCoE ports.
- ► Connectivity information is shown between CNAs and FCoE ports on the switch within the Topology Viewer.
- ► Tivoli Storage Productivity Center displays both FCoE and FC ports in the switch ports list. For the Brocade FCoE switch, domain-port zoning is limited to the FC ports and only worldwide name (WWN) zoning is supported on the FCoE ports. For the Cisco Nexus switches, domain-port zoning is not supported.

14.2 Additional switch models supported

In Tivoli Storage Productivity Center V4.2, many new switches are supported. Included in this list are FCoE capable switches, which are currently supported on a tolerance level basis. This means that extensive testing was performed to ensure sure that their presence in your SAN environment will not affect any Tivoli Storage Productivity Center functions, even if FCoE and CEE capabilities are still not supported.

14.2.1 Brocade 3016, 5410, 5470, 5480, 7800, and M5440

All switch functions are supported.

14.2.2 Brocade Silk Worm 7800 (IBM 2005-R06)

Only switch performance monitoring is supported.

14.2.3 Brocade DCX-4S Backbone

Brocade DCX-4S Backbone is a smaller model of Brocade DCX (a larger model) and is supported with previous Tivoli Storage Productivity Center releases. All switch functions are supported.

14.2.4 Brocade 8000

The Brocade 8000 is a top-of-rack link layer (Layer 2) CEE/FCoE switch with 24 10 Gigabit Ethernet (GbE) ports for LAN connections and eight Fibre Channel ports (with up to 8 Gbps speed) for Fibre Channel SAN connections

The Brocade 8000 support limitations are as follows:

- In Tivoli Storage Productivity Center V4.2, only tolerance level support for Brocade 8000 switches is added. CEE and FCoE capabilities are not supported. However, all the FC and FCoE ports are displayed in Tivoli Storage Productivity Center.
- ▶ Domain-port zoning in Brocade 8000 is limited to the eight FC ports. For the internal FCoE ports, only WWN zoning is supported.
- ► As with prior releases, use SMIS agent for Brocade FCoE switches rather than using Storage Resource agent or out-of-band Fabric agent.
- Collecting performance data and zone control functions can be done using only the SMIS agent.

14.2.5 Cisco Nexus 5000 Series

The Cisco Nexus 5000 Series includes the Cisco Nexus 5020 switch and the Cisco Nexus 5010 switch.

Fibre Channel ports are optional on the Cisco Nexus 5000 Series switch. The expansion modules up to 8 Fibre Channel ports are available on the Cisco Nexus 5010 switch and up to 16 Fibre Channel ports are available on the Cisco Nexus 5020 switch. Tivoli Storage Productivity Center can discover the switch even if it does not have any FC ports.

The Cisco Nexus 5000 Series support limitations are as follows:

- In Tivoli Storage Productivity Center V4.2. tolerance level support has been added for Cisco Nexus 5000 switches. Currently, Tivoli Storage Productivity Center does not support the CEE/FCoE capabilities of these switches.
- Tivoli Storage Productivity Center V4.2 displays Fibre Channel and virtual Fibre Channel ports. The Fibre Channel portion of FCoE is configured as a virtual Fibre Channel interface.
- As with Cisco MDS switches, the domain-port zoning is not supported, and zone control functions can only be done using the Storage Resource agent.

14.3 Additional HBA and CNA models supported

Tivoli Storage Productivity Center V4.2 adds support for the following HBA and CNA models:

- 4 Gbps HBAs:
 - Brocade 415 / 425
 - Emulex LP1105-BC (Blade Servers)
 - HP AD300A (HP-UX Itanium)
 - QLogic QLA2440, QLE2360 / 2362, QLE2440
- ▶ 8 Gbps HBAs:
 - Brocade 815 / 825
 - QLogic QLE2564 (Standard Servers), QMI2742 (Blade Servers)
- 10 GbpsE CNAs (toleration mode support):
 - Brocade 1010 / 1020
 - Emulex LP21000 / 21002
 - QLogic QLE8150 / 8152, QLE8140 / 8142, QLE8042

14.4 Integration with Brocade Data Center Fabric Manager and Brocade Network Advisor

Brocade Distributed Data Center Fabric Manager (DCFM) is a central management solution for Brocade directors, routers, and switches, HBAs, and CNAs. DCFM integrates with Tivoli Storage Productivity Center V4.2 by adding the DCFM server as a CIMOM Agent.

Brocade Network Advisor is a central network management solution for SAN and IP networks. It replaces Brocade DCFM and integrates similarly with Tivoli Storage Productivity Center for adding the CIMOM Agent.

Both of these applications include an integrated SMI Agent that can manage both Brocade b-type (FOS) and m-type (formerly McDATA EOS) fabrics. A separate proxy SMI agent is no longer required.

Functions: Using Network Advisor as a CIMOM Agent is supported starting with Tivoli Storage Productivity Center 4.2.1 fix pack 4. Although Network Advisor can manage IP networks, that information is not collected by Tivoli Storage Productivity Center. No additional functionality is enabled by using Network Advisor over one of the other SMIS Agents.

14.4.1 Supported functions

Tivoli Storage Productivity Center can access the fabrics managed by DCFM or Network Advisor, supporting the following functions:

- Fabric Topology and Zone Probe
- Switch Performance Monitoring
- ► Zone Control:
 - Empty zone sets, zones, and aliases are not supported
 - Zone aliases are only supported for b-type fabrics
- ► Alerts

Versions: The minimum versions qualified with Tivoli Storage Productivity Center are as follows:

- ► DCFM 10.4.1
- ► Network Advisor 11.1

For additional support information, refer to the Supported Products Matrix at this website:

https://www-304.ibm.com/support/docview.wss?uid=swg21386446

Always refer to the Brocade documentation to ensure that the DCFM or Network Advisor version you are using supports the switch models and firmware levels in your environment.

14.4.2 Configuring the integrated SMI agent for DCFM and Network Advisor

To use the integrated SMI Agent that is a part of DCFM and Network Advisor as a CIMOM Agent with Tivoli Storage Productivity Center, you must have it installed and configured as part of the DCFM or Network Advisor application.

If you did not configure the SMI Agent at installation time, or if you want to change some of the properties, you can do this by launching the **Brocade Server Management Console** and then clicking **Configure SMI Agent...** to launch the configuration panel.

Further details for installing and configuring DCFM or Network Advisor and its integrated SMI Agent can be found in the following Brocade publications that are available on the MyBrocade portal of the Brocade website:

- DCFM Installation Guide
- ► DCFM Enterprise User Manual
- ► Brocade Network Advisor Installation Guide
- Brocade Network Advisor SAN User Manual

14.4.3 Adding a DCFM or Network Advisor server into Tivoli Storage Productivity Center

Complete the following steps:

1. Click Configure Devices in the CIMOM Agents panel (Figure 14-1).

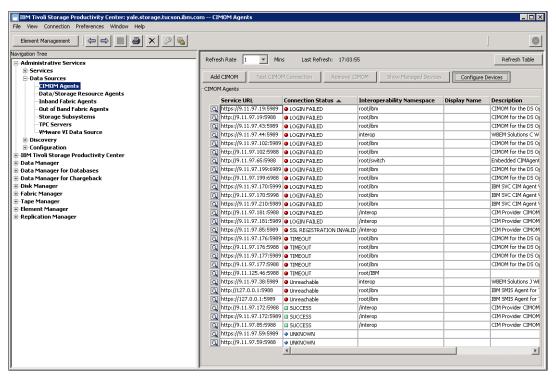


Figure 14-1 CIMOM Agents panel

2. Select Fabrics/Switches (Figure 14-2) and click Next to continue.

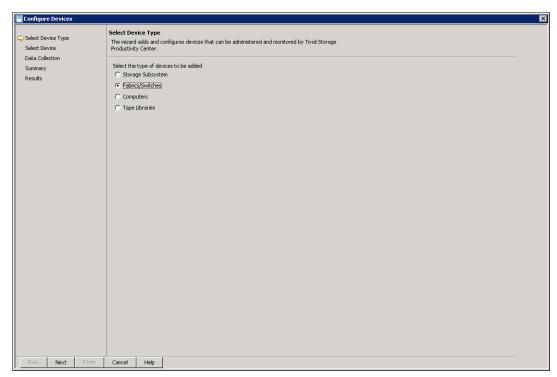


Figure 14-2 Select device type

3. Select Add and configure new fabrics (Figure 14-3), and click Next to continue.

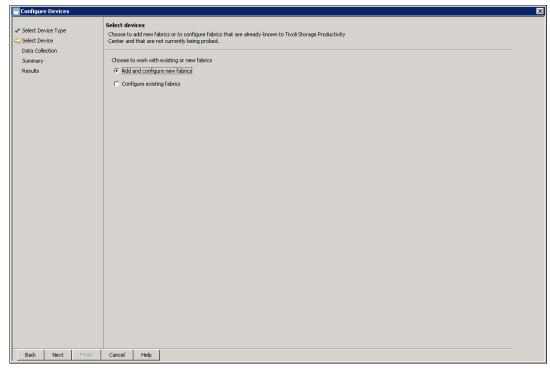


Figure 14-3 Select devices

4. Select **Configure a CIMOM Agent for monitoring the fabric**, fill in your DCFM or Network Advisor server information (Figure 14-4), and then click **Add**.

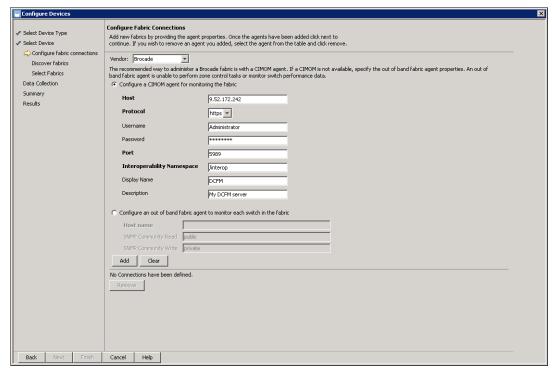


Figure 14-4 Configure fabric connections

Note: For a DCFM or Network Advisor CIMOM, the default values are as follows:

► Protocol: HTTPS

► User name: Administrator

► Password: password

► Interoperability Namespace: /interop

Notice the uppercase "A" in the username.

 The DCFM or Network Advisor server is added to the list at the bottom of the panel (Figure 14-5). You can add more DCFM or Network Advisor servers, or click Next to continue.

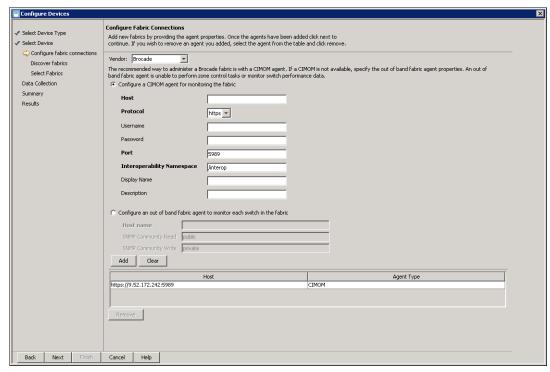


Figure 14-5 DCFM server added

A Fabric discovery job is initiated (Figure 14-6).

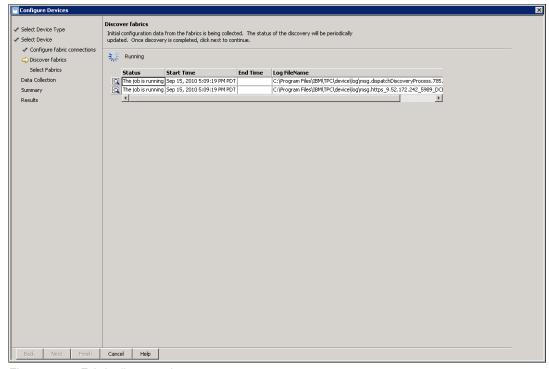


Figure 14-6 Fabric discovery in process

6. When the Fabric discovery is done, click **Next** to continue (Figure 14-7).

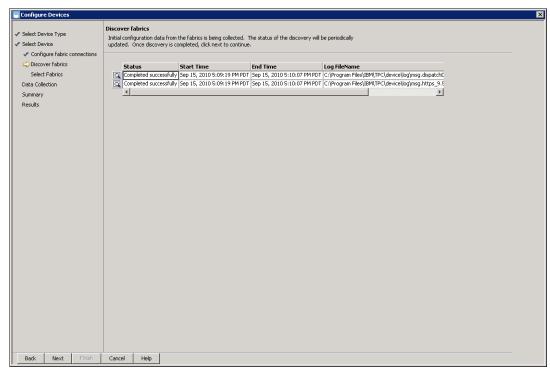


Figure 14-7 Discovery completed

7. A list with every newly discovered fabric is displayed (Figure 14-8). These are the same fabrics under DCFM or Network Advisor management. Select the ones you intend to manage with Tivoli Storage Productivity Center, or all fabrics, and click **Next** to continue.

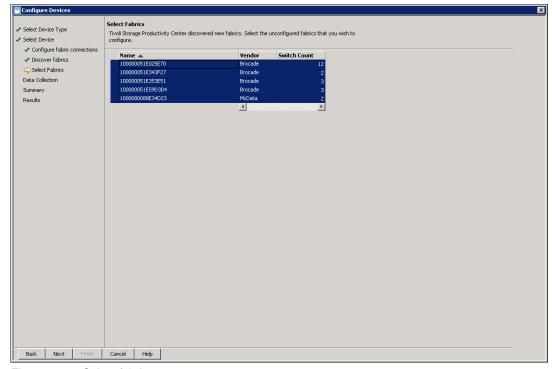


Figure 14-8 Select fabrics

8. Use the next panel (Figure 14-9) to add the newly discovered fabrics to a previously defined Monitoring Group. Select a group and click **Next** to continue.

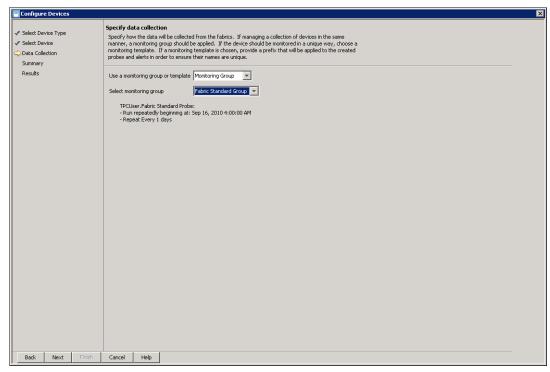


Figure 14-9 Fabric Monitoring Groups

9. Review the summary of your choices (Figure 14-10). Click Next to continue.

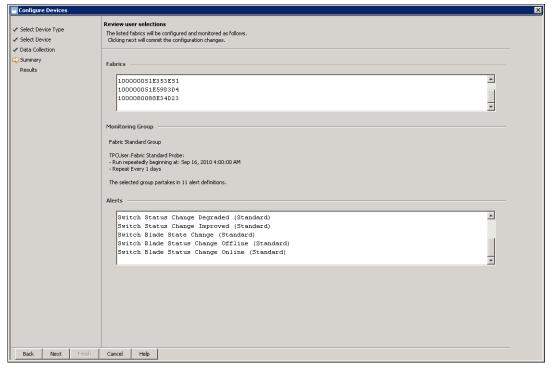


Figure 14-10 User selections review

10. Tivoli Storage Productivity Center processes the changes and displays the results (Figure 14-11). You have successfully added your DCFM or Network Advisor managed fabrics to your Tivoli Storage Productivity Center environment.

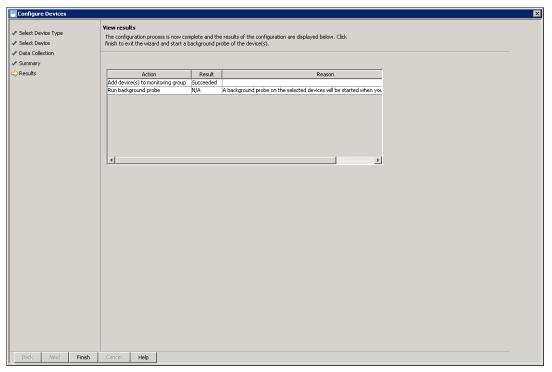


Figure 14-11 Results panel



15

Reporting

In this chapter, we focus on the built-in reporting functions, including the enhancements for Storwize V7000 and SVC internal (local) disk reports of Tivoli Storage Productivity Center V4.2. We also introduce the new Storage Tier reports, and provide useful reporting scenarios.

Storage Tier reports help you decide where to place frequently and infrequently accessed data to get the best performance from your system. Infrequently accessed data can be placed on lower-cost storage systems and frequently accessed data can be placed on faster performing storage systems. We show you how to produce these reports and provide overall directions on how to use them.

15.1 Reporting overview

Tivoli Storage Productivity Center collects a wide range of information from the storage environment it is monitoring. All collected data, including configuration data, utilization data, and performance data is stored in a relational DB2 database schema. This chapter provides an overview that illustrates the various reports that Tivoli Storage Productivity Center can generate from this collected data.

Tip: The Reporting menu structure has not changed since Tivoli Storage Productivity Center V4.1.

Before you can use reports to view details about your environment, you must collect information from it. You can use monitoring jobs in Tivoli Storage Productivity Center such as probes, scans, pings, and Tivoli Storage Productivity Center server probes, and discovery jobs to gather comprehensive information and statistics about your storage resources.

These are the multiple options available to extract data from Tivoli Storage Productivity Center.

15.1.1 Tivoli Storage Productivity Center GUI

The Tivoli Storage Productivity Center GUI provides access to a large number of predefined reports, provides charting and graphing functions, allows customizing reports (up to a certain extent) and scheduling them to run on a repeated basis (called *Batch Reports*). In addition, a graphical Topology Viewer is provided so you can interactively browse the current, and historic states of the monitored storage environment (Configuration History). These reports can also be displayed with two types of charts: graphical charts or as tables of data.

15.1.2 Batch reports

The Tivoli Storage Productivity Center GUI allows definition of Batch Reports, which are typically scheduled to run repeatedly, producing output in either comma separated values (CSV), plain text, PDF, or HTML format. A subset of the reports available through the GUI can be exported that way to allow for automatic publishing or post-processing of Tivoli Storage Productivity Center data, without having to invoke either GUI or CLI. Batch reports might require installation of a Tivoli Storage Productivity Center Data agent to run, or the Storage Resource agent. They are available only with the Tivoli Storage Productivity Center for Data or Tivoli Storage Productivity Center Standard Edition licenses.

15.1.3 Tivoli Storage Productivity Center CLI (TPCTOOL)

The Tivoli Storage Productivity Center command-line interface (CLI) is also known as TPCTOOL. It provides programmatic access to a subset of the reports that are available through the GUI, primarily performance reports, and a limited number of configuration reports. For additional information about reporting by Tivoli Storage Productivity Center CLI (and TPCTOOL as a reporting tool), see *SAN Storage Performance Management Using Tivoli Storage Productivity Center*, SG24-7364. In addition, see *Reporting with TPCTOOL*, REDP-4230 for detailed information about how to use the Tivoli Storage Productivity Center CLI to generate reports.

15.1.4 Reporting database views

Introduced in Tivoli Storage Productivity Center V4.1, a set of database views is available to access data from DB2 directly using Structured Query Language (SQL). The views represent the raw data "behind" the reports that can be produced through GUI and CLI, thus allowing flexible customizations. Notice that the database views do not allow exporting the reports readily available through GUI and CLI directly, but instead provide access to the actual data within the database repository. When generating reports through Tivoli Storage Productivity Center GUI, for example, additional calculations are performed on the data in the repository. SQL access to database views are covered in detail in the publication *IBM Tivoli Storage Productivity Center V4.1 Release Guide*, SG24-7725 in the section titled "Database repository access".

15.2 Generating reports

Use Tivoli Storage Productivity Center reports to view detailed information about the storage resources in your environment. Reports in the Tivoli Storage Productivity Center GUI are organized into reporting types and reporting categories:

- ► Reporting types define how the report is handled by the system (run online or offline, user-defined or system-defined. Types are as follows:
 - System reports
 - Batch reports
 - User ID reports
 - User-defined online reports
- Reporting categories are intended to organize the vast number of available reports and help users find the information that they are particularly interested in. Categories are as follows:
 - Asset, availability
 - Capacity
 - Usage
 - Usage violation
 - Backup
 - Storage subsystem performance
 - Switch performance

For detail information about the various reporting types, see *IBM Tivoli Storage Productivity Center Version 4.2 User's Guide*, SC27-2338, chapter 6.

15.3 Reporting example

In this section, we introduce the Tivoli Storage Productivity Center reporting structure. Also, we show how to generate Tivoli Storage Productivity Center reports from the GUI, TPCREPORT with SQL statements, TPCTOOL and BIRT.

First, we show the Tivoli Storage Productivity Center Navigation Tree, which has several subtrees. For most of the manager subtrees, you are able to generate various kinds of reports. This structure is shown in Figure 15-1.

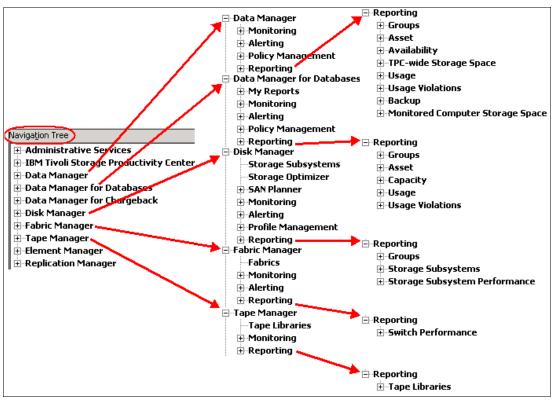


Figure 15-1 Tivoli Storage Productivity Center reporting structure

Each subtree allows you to generate reports depending on what information interests you. Next, through an example, we show the steps for report generation and the options available.

1. Choose the report type.

In our example, we consider the Storage Subsystem Disks report. From the Navigation Tree, go to **Disk Manager** \rightarrow **Reporting** \rightarrow **Storage Subsystems** \rightarrow **Disks** \rightarrow **By Storage Subsystem**, as shown in Figure 15-2.

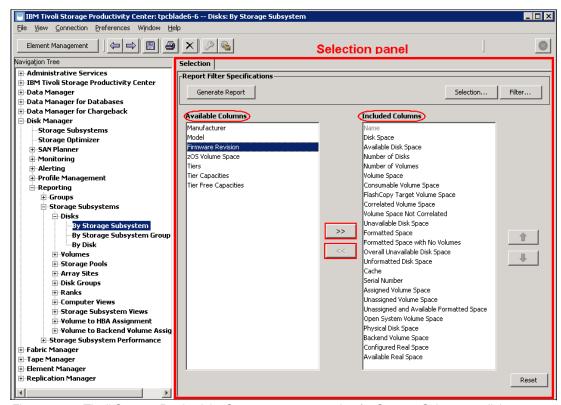


Figure 15-2 Tivoli Storage Productivity Center report generation for Storage Subsystem disks

2. Select the columns.

From the Available Columns in the Selection panel (Figure 15-2), you can select the columns to add to your report. The Included Columns section shows the columns that are already included and will appear in the generated report. Use the arrows to add or remove columns.

3. Select resources to include.

Click **Selection** to open the Select Resources window (Figure 15-3). Use this window to select, for example, the Storage Subsystems you want to include in your report. Click **OK** to accept the changes.

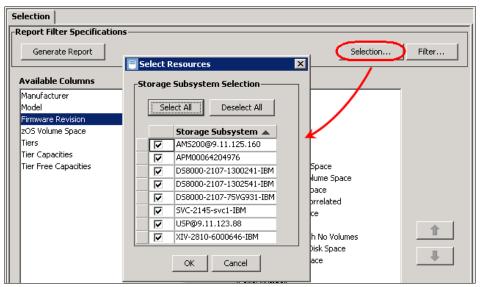


Figure 15-3 Tivoli Storage Productivity Center report, Select Resources window

4. Add filters.

Click **Filter** to open the Edit Filter window (Figure 15-4), which is where you can add filters. You can create filters that allow you to create conditions on column values using operators such as equal, greater than, distinct, like, between, and null. You can select that all the records meet the filters you define or that at least one condition is met. Click **OK** when you finish.

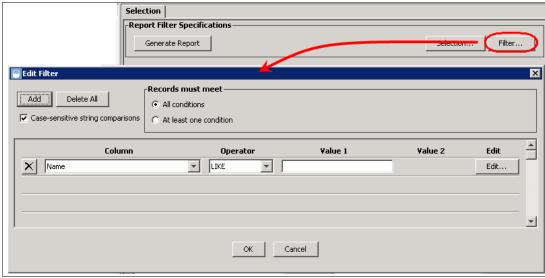


Figure 15-4 Tivoli Storage Productivity Center report, Edit Filter window

5. Generate the report.

Now you are ready to generate the report. Click **Generate Report** to open the Storage Subsystems panel that shows the output generated (Figure 15-5). Note that not all the columns are included in the graphic shown. The report also shows, in the yellow row, the consolidated values for the columns included previously.

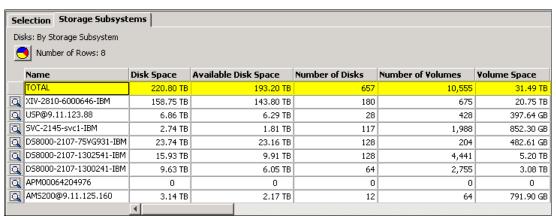


Figure 15-5 Generated report

From the panel shown in Figure 15-5, you can click the tri-color button () to generate various charts for this report. As shown in Figure 15-6, certain options become disabled, based on the report that is generated.



Figure 15-6 Chart selection

If the menu options are enabled, however, click **Chart: All Rows** to open the panel shown in Figure 15-7. The report now is presented in a Column (Bar) chart.

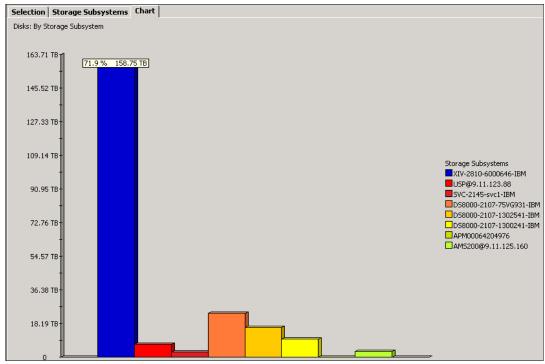


Figure 15-7 Column chart for the generated report

In Figure 15-5 on page 573, you can highlight rows from the output and generate charts based on your selection. In this way, you can view historical data for the selected resources.

Next, we consider a performance report that shows data suited for historical charts.

For this new report, navigate to Disk Manager → Reporting → Storage Subsystem
 Performance → By Storage Subsystem. The panel is similar to the one in Figure 15-8.

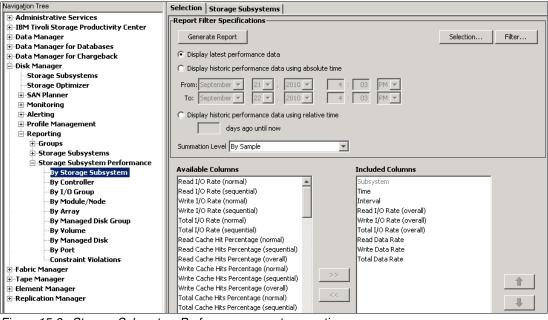


Figure 15-8 Storage Subsystem Performance report generation

- 2. To produce a performance report, select the columns to be included, edit filters, and select the storage subsystems by clicking **Selection**. Select the timeframe for the report and the summation level as shown in Figure 15-9.
 - Options for the time frame are as follows:
 - Select the latest data.
 - Specify the starting and ending date.
 - · Use relative time.

Options for Summation Level are as follows:

- By Sample means that the report will contain a data point for each performance data sample gathered.
- Hourly means that calculates the average during one hour intervals.
- Daily is for the daily average.

Click Generate Report when you are done with your selections.

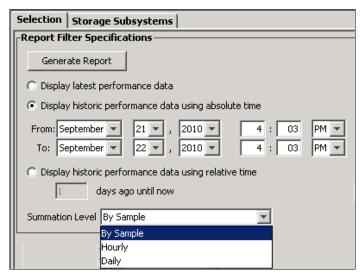


Figure 15-9 Performance report specification

After the performance report is generated, output similar to Figure 15-10 is displayed.

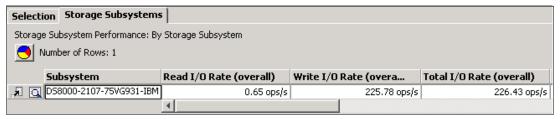


Figure 15-10 Performance report

3. Click the tri-color button () to see a more detailed selection panel for the charts. This panel (Figure 15-11) is the Select Charting Options panel.

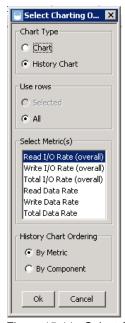


Figure 15-11 Select Charting Options panel

- a. Select the Chart Type:
 - Chart opens a chart with only one measure.
 - History Chart generates historical data.
- b. The Use Rows section is where you select specific rows from the generated report.
- c. The Select Metrics section generates the chart based on one of the "Included Columns" of the report.
- d. The History Chart Ordering section is where you order the chart in either of the following ways:
 - By Metric (read I/O rates, and so on)
 - By Component (volumes, ports, and so on)
- e. Click **OK** to generate the report based on your selections.

For our environment, the report looks like Figure 15-12. From the generated report you are able to change the time interval and also the aggregation of samples to By sample, Hourly, and Daily.

From the report, you also have the option to show the future trends based on the historical information Tivoli Storage Productivity Center has gathered. Clicking **Show Trends** shows a dashed line that follows the solid one, as shown in Figure 15-12.

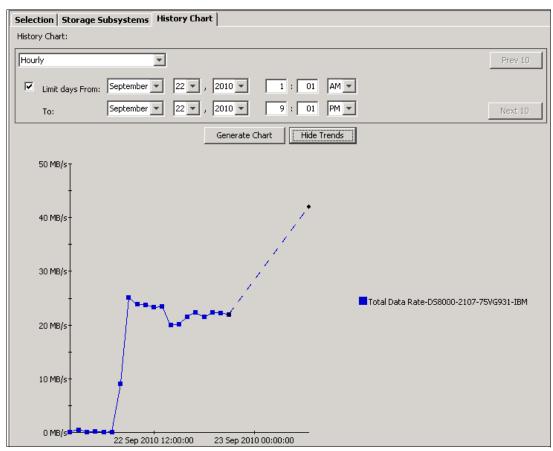


Figure 15-12 Storage Subsystem Performance report showing trends

We have presented the steps needed to generate a report through the GUI. You can generate the same report through the TPCREPORT schema using SQL. For more details, see the chapter about reporting in the *IBM Tivoli Storage Productivity Center V4.1 Release Guide*, SG24-7725.

15.4 Storage Capacity Pyramid

The Storage Capacity Pyramid is a result of customers who diligently monitored their installed storage and utilized storage. These customers invariably had many Excel spreadsheets with information about storage system capacity and utilization, without an understanding of how to find what they needed within the Tivoli Storage Productivity Center.

15.4.1 Overview

The Storage Capacity Pyramid (Figure 15-13) was put together to show true capacity and where space "leaks" out of total capacity at many points within a configuration to end up with usable capacity. So, a goal is to understand how storage is utilized, starting from raw drives to file system capacity. Keep in mind that this is not an attempt to cover all situations, and depending on your environment, you might require additions to the pyramid.

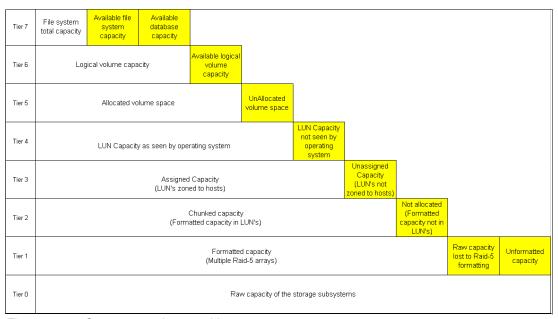


Figure 15-13 Storage capacity pyramid

15.4.2 Tiers

We examine the tiers from the bottom up, showing how to obtain Tivoli Storage Productivity Center reports for each one. We also explain the Tivoli Storage Productivity Center definition for each report with considerations to take into count. Within each definition, between parentheses, is the original Tivoli Storage Productivity Center name of the column, because some names have been updated. Also, each concept in the tier is matched, by using numbers in the red circles, to concepts on the Tivoli Storage Productivity Center reports, for example, as shown in Figure 15-14.

Tier 0: Raw capacity of the storage subsystems

The definition is shown in Figure 15-14.

Raw capacity of the storage subsystems

Figure 15-14 Tier 0 definition

The raw capacity of a storage subsystem is based on the Tivoli Storage Productivity Center report on Disk Space, which is defined in Table 15-1.

Table 15-1 Disk Space definition

Disk Space (Disk Capacity)

Total raw (unformatted) disk capacity of a storage subsystem. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

Capacity of spare disks identified on storage subsystems.

Capacity of storage subsystem disks that become missing after a Tivoli Storage Productivity Center storage subsystem probe.

The value in the TOTAL row for this column represents the sum of the capacity on disks detected by Tivoli Storage Productivity Center. For SAN Volume Controller, the value in the TOTAL row represents the sum of the capacity of MDisks detected by Tivoli Storage Productivity Center.

To see this report, go to **Data Manager** \rightarrow **Reporting** \rightarrow **Asset** \rightarrow **System-wide** \rightarrow **Storage Subsystems**. For our environment, the report is similar to Figure 15-15.

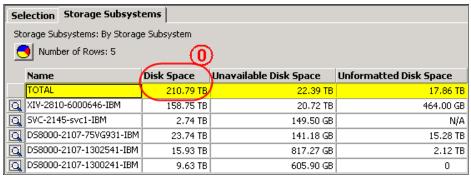


Figure 15-15 Tivoli Storage Productivity Center raw physical capacity report

Tier 1: Formatted, raw capacity lost to RAID 5, and unformatted capacity

Tier 1 is formed by formatted capacity, raw capacity lost to RAID 5 formatting and unformatted capacity. See Figure 15-16.



Figure 15-16 Tier 1 definition

These concepts match Tivoli Storage Productivity Center concepts shown in Table 15-2.

Table 15-2 Formatted Space, Unavailable Disk Space, and Unformatted Disk Space

Formatted Space (Formatted Capacity)	Total amount of formatted storage space that is associated with the storage subsystem that is used or can be used for subsystem storage volumes or virtual storage volumes. This value depends on the storage subsystem SMI-S provider returning a formatted value. For certain storage subsystems, this space that cannot be used for storage volumes or virtual storage volumes is the total managed space, which includes space that cannot be used for storage volumes or virtual storage volumes. This value can occur for various reasons, such as space not being formatted until storage volumes are created. For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.
Unavailable Disk Space (Overhead)	Amount of storage subsystem volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage subsystem SMI-S provider returning a valid RAID value that is used to determine the overhead. If the value is zero for a storage subsystem, the overhead cannot be calculated. This amount does not include storage space information from storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
Unformatted Disk Space (Undefined Disk Group or Array Site Capacity)	This value does not apply to all storage subsystems. For IBM FAStT, IBM Enterprise Storage Server, IBM DS6000, IBM XIV, and IBM DS8000 storage subsystems, this value is the total raw disk space of any unformatted disk groups or array sites. For IBM XIV Storage System, this value is the total available physical (hard) space that is left for pool creation. This column displays N/A for the storage subsystems where this value is not applicable.

To get this report, go to **Data Manager** \rightarrow **Reporting** \rightarrow **Asset** \rightarrow **System-wide** \rightarrow **Storage Subsystem**. Make sure that you include the columns previously defined. The generated output is shown in Figure 15-17.

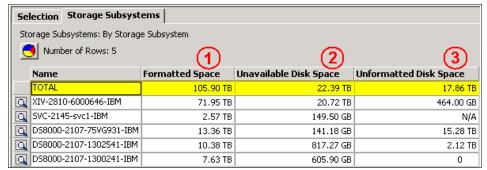


Figure 15-17 Formatted space, Unavailable disk space and Unformatted disk space report

Tier 2: Chunked capacity and not allocated capacity

For Tier 2, the chunked capacity and not allocated capacity are considered. See Figure 15-18.



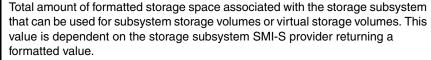
Figure 15-18 Tier 2 definition

The Tivoli Storage Productivity Center report columns that match these are detailed in the Table 15-3.

Table 15-3 Formatted Space, Unavailable Disk Space, and Unformatted Disk Space

Consumable Total amount of unique storage subsystem volume space and virtual storage Volume Space volume space on monitored storage subsystem arrays that you can assign to (Usable LUN servers within the network or are already assigned to servers within the network. Capacity) Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column: Storage subsystem volumes or virtual storage volumes that become missing after a probe. Storage subsystem volumes or virtual storage volumes that are used as flash copy targets. Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx). Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database. Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers. For example, IBM San Volume Controller. Note: For IBM FAStT, HP StorageWorks, and Hitachi Data Systems storage subsystems, the value in this column is the same as the value in the Volume Space column.

Formatted Space with No Volumes (Formatted Space with No Volumes)





For some subsystems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur because of a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the space within the managed disk groups that does not include the size of the virtual disks created and the corresponding virtual disk extent overhead.

For report generation go to **Data Manager** \rightarrow **Reporting** \rightarrow **Asset** \rightarrow **System-wide** \rightarrow **Storage Subsystems** \rightarrow **By Storage Subsystem**. The resulting output is shown in Figure 15-19.

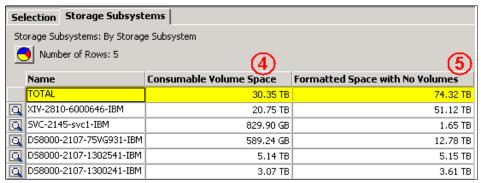


Figure 15-19 Consumable volume space and Formatted space with no volumes report

Tier 3: Assigned and unassigned capacity

Assigned capacity and unassigned capacity are considered for Tier 3. See Figure 15-20.

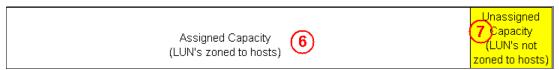


Figure 15-20 Tier 3 definition

The Tivoli Storage Productivity Center report columns that match these are detailed in Table 15-4.

Table 15-4 Assigned Volume Space and Unassigned Volume Space

Assigned Volume
Space
(Subsystem
Assigned LUN
Capacity)

6

Total storage subsystem volume space within the storage subsystem that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.

Storage subsystem volumes that are not mapped or are not assigned to host systems.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)



Total storage subsystem volume space within a storage subsystem that is not mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- ▶ Storage subsystem volumes that are mapped or assigned to host systems.

Go to **Data Manager** → **Reporting** → **Asset** → **System-wide** → **Storage Subsystems** → **By Storage Subsystem**. The resulting output for the generated report is shown in Figure 15-21.

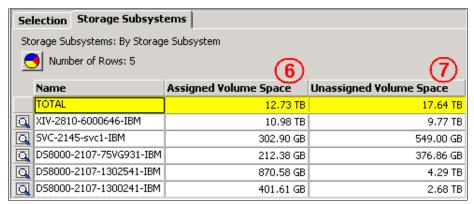


Figure 15-21 Tivoli Storage Productivity Center report on Assigned and Unassigned volume space

Tier 4: LUN capacity as seen and not seen by operating system

LUN capacity as seen by the operating system, and LUN capacity not seen by the operating system are considered for Tier 4. See Figure 15-22.

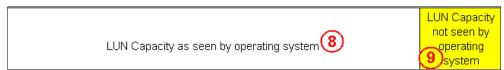


Figure 15-22 Tier 4 definition

Tier 4 is LUN Capacity seen and not seen by operating systems. Table 15-5 shows the Tivoli Storage Productivity Center report columns for these. Take into account that, to get these reports, you go to various menus described next.

Table 15-5 Volume Space from computer views and not visible to monitored servers menus

Volume Space (Volume Capacity)	Storage capacity that is on a volume
Volume Space (Volume Capacity)	Storage capacity that is on a volume. Note: For space-efficient volumes and space-efficient VDisks, the volume or VDisk capacity is the amount of storage space that is requested for these volumes or VDisks, not the actual allocated amount. The result is in discrepancies in the overall storage space that is reported for a storage subsystem using space-efficient volumes or VDisks. This also applies to other space calculations, such as the Consumable Volume Space and FlashCopy Target Volume Space calculations.

To generate a report containing "Volume space" from the Computer Views menu, go to **Disk Manager** \rightarrow **Reporting** \rightarrow **Storage Subsystems** \rightarrow **Computer Views** \rightarrow **By Computer**. The output is shown in Figure 15-23.

Selection C	omputers		
Storage Subsystem: By Computer			
Number of Rows: 5			8)
Computer	Device Path (computer)	Storage Subsystem	Volume Space
TOTAL			5,00 GB
TPCBLADE4-1	3 Disk 33	SVC-2145-svc1-IBM	1.00 GB
TPCBLADE4-1	3 Disk 32	SVC-2145-svc1-IBM	1.00 GB
TPCBLADE4-1	3 Disk 27	SVC-2145-svc1-IBM	1.00 GB
TPCBLADE4-1	3 Disk 26	SVC-2145-svc1-IBM	1.00 GB
TPCBLADE4-1	3 Disk 18	SVC-2145-svc1-IBM	1.00 GB

Figure 15-23 Tivoli Storage Productivity Center Volume Space column

To generate a report containing "Volume space" from the Not Visible to Monitored Server menu, go to **Disk Manager** \rightarrow **Reporting** \rightarrow **Storage Subsystems** \rightarrow **Volume to HBA Assignment** \rightarrow **Not Visible to Monitored Server**. The generated output is shown in Figure 15-24.

This information is available only if Data agents are installed on the host computers. Configuration information that is collected by the Data agent is then correlated with storage subsystem information that is collected by the CIMOM.

Sel	ection Storage Subsyst	ems		
Volume HBA Assignment: Not Visible to Monitored Computers Number of Rows: 88				
	Storage Subsystem	Volume Name	Volume Space	Computer
	TOTAL		165.81 GB	
Q	SVC-2145-svc1-IBM	svtest1	2.00 GB	TPCBLADE4-11
	SVC-2145-svc1-IBM	svttest	4.00 GB	TPCBLADE4-11
	SVC-2145-svc1-IBM	blade411	2.00 GB	TPCBLADE4-11
	SVC-2145-svc1-IBM	blade411	2.00 GB	TPCBLADE4-11
ℚ	SVC-2145-svc1-IBM	dtest1	15.00 GB	TPCBLADE4-11

Figure 15-24 "Volume space" column, "Not visible to monitored server" report

Tier 5: Allocated and unallocated volume space

Tier 5 consists of allocated and unallocated volume space. See Figure 15-25.



Figure 15-25 Tier 5 definition

Table 15-6 shows the Tivoli Storage Productivity Center report columns for these.

Table 15-6 Consumed Volume Space and Available Volume Space columns

Consumed Volume Space (Allocated Volume Space)	Storage capacity of a volume minus unallocated space.
Available Volume Space (Unallocated Volume Space)	Volume capacity that is assigned to each computer minus the space that has been allocated to a logical volume.

To generate this report navigate to **Disk Manager** \rightarrow **Reporting** \rightarrow **Storage Subsystems** \rightarrow **Computer views** \rightarrow **By Computer**. The generated output is shown in Figure 15-26.

Note that this information is only available if Data agents exist on the host computers. The Data agent-collected configuration information is then correlated with the CIMOM-collected Storage System information.

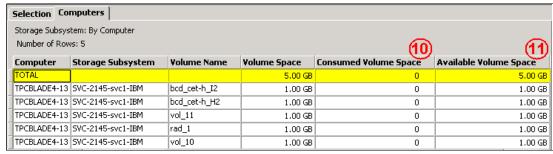


Figure 15-26 Consumed and Available Volume Space reports

Tier 6: Logical volume capacity and available logical volume capacity

Tier 6 consists of logical volume capacity and available logical volume capacity. See Figure 15-27.

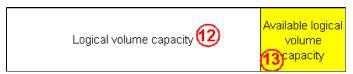


Figure 15-27 Tier 6 definition

These concepts are described in Table 15-7.

Table 15-7 Consumed Volume Space and Available Volume Space columns

Consumed Volume Space (Allocated Volume Space)	Each row in the report gives you detailed information for each disk or volume group. The following information is displayed: Computer Disk/Volume Group Name (Disk/Volume) group name Disk/Volume Group Space (Capacity) 12 Available Disk/Volume Group Space (Freespace) 13 Percent Available Disk/Volume Group Space Type Number of Disks Number of Logical Volumes or Filesystems (Number of file systems/logical volumes) Last Probe Time Discovered Time
---	---

Go to Data Manager \rightarrow Reporting \rightarrow Asset \rightarrow System-wide \rightarrow Disk/Volume Groups \rightarrow By computer. Figure 15-28 shows the generated report.

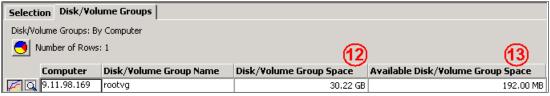


Figure 15-28 Disk/Volume Group Space" and "Available Disk/Volume Group Space" report

Tier 7: File system total, and available file system and database capacity

For Tier 7, File system total capacity, available file system capacity and available database capacity are considered. See Figure 15-29.



Figure 15-29 Tier 7 definition

These concepts are described in Table 15-8. First, we show how to generate a report on the first two concepts and next how to generate it for Available database capacity.

Table 15-8 File System Space, Available File System Space, and File Free Space columns

File System Space (Capacity)	 Amount of file system storage space for a file system. Notes: ► The value that appears in the TOTAL row for this column does not include file systems on virtual machines where the Data agent is installed. ► The value in this column might be greater than Used Space plust Free Space on some UNIX systems. This value occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the Capacity field is greater than or equal to the value of Used Space and Free Space, the difference representing the space wasted because of the system overhead.
Available File System Space (Free Space)	Amount of unused storage space within a file system.
File Free Space (Container Free Space)	Amount of free space available to the Instances on a computer.

To generate a report on Total File System Capacity and Free Space, go to **Data Manager** \rightarrow **Reporting** \rightarrow **TPC-wide Storage Space** \rightarrow **File System Space** \rightarrow **By Filesystem**. Its output is shown in Figure 15-30. This information is available only if there are Storage Resource agent or Data agents on the host computers.

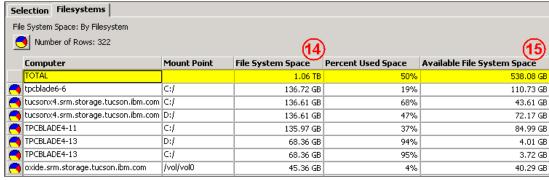


Figure 15-30 File System Space" and "Available File System Space" report

To generate a report on Available Database Capacity, go to **Data Manager for Databases** \rightarrow **Reporting** \rightarrow **Capacity** \rightarrow **All DBMSs** \rightarrow **Total Instance Storage** \rightarrow **By Computer**. The resulting output is shown in Figure 15-31.

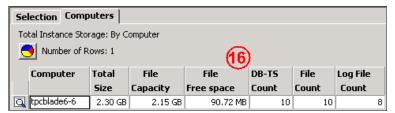


Figure 15-31 Tivoli Storage Productivity Center report on File Free Space

15.5 Capacity licensing

In this section we explain how to determine the licensed capacity you are currently using in your environment. Before following the steps needed to determine the capacity that is monitored by Tivoli Storage Productivity Center and the scenarios covered, we explain several necessary concepts:

Formatted capacity:

Formatted capacity is the total amount of formatted storage space associated with the storage system that is used or can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system returning a formatted value.

► Monitored computer space:

Use these reports to view information about disks that reside on computers that are monitored by IBM Tivoli Storage Productivity Center, including information about storage subsystem disks, known disks with serial numbers, or unknown disks without serial numbers.

For more information about monitored computer space, the Redpaper publication, *IBM Tivoli Storage Productivity Center V4.2, Monitored Computer Storage Space Reports*, REDP-4820.

http://www.ibm.com/redbooks

Correlated definition:

Indicates *Yes* if Tivoli Storage Productivity Center is able to relate the host storage directly to a storage subsystem.

We cover the following scenarios using separate versions of Tivoli Storage Productivity Center:

- Tivoli Storage Productivity Center for Disk based on "Formatted capacity"
- Tivoli Storage Productivity Center for Data based on "Monitored computer space"
- Tivoli Storage Productivity Center Standard Edition combining Disk and Data without overlap

15.5.1 Tivoli Storage Productivity Center for Disk Licensing

Tivoli Storage Productivity Center for Disk license is based on the amount of Formatted Capacity, which is the total amount of space currently visible by connection to storage subsystem minus the virtualized storage already counted as back-end storage. It does not double count storage that is virtualized behind SVC or Storwize V7000.

The following procedure gets the total Formatted Space for your environment:

1. Go to Disk Manager → Reporting → Storage Subsystems → Disks → By Storage Subsystem as shown in Figure 15-32.

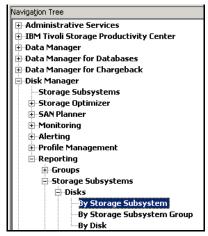


Figure 15-32 By Storage Subsystem

2. From the Selection panel, make sure to include only the columns shown in Figure 15-33.

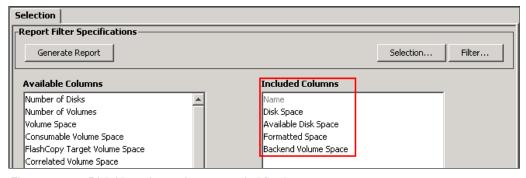


Figure 15-33 Disk Licensing, columns needed for the report

 From the Selection panel, click Selection. The Storage Subsystem Selection panel opens (Figure 15-34). Identify your SVC Storage Subsystems and deselect them. Click OK. Back on the Selection panel, click Generate Report.

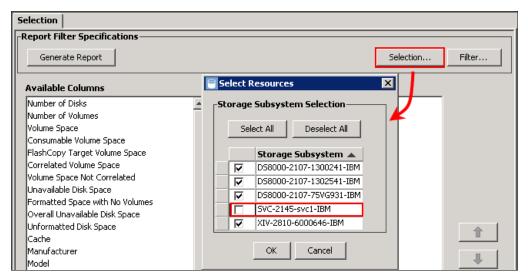


Figure 15-34 Disk Licensing, Storage Subsystem Selection panel

4. Review the resulting report, which shows the formatted space that Tivoli Storage Productivity Center is currently managing. For example, Figure 15-35 shows that Tivoli Storage Productivity Center is managing 103.33 TB.

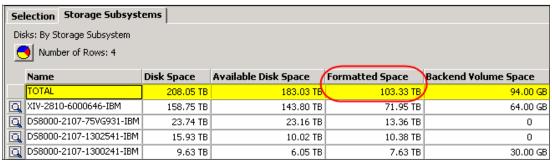


Figure 15-35 Disk Licensing, report generated

15.5.2 Tivoli Storage Productivity Center for Data Licensing

Tivoli Storage Productivity Center for Data licensing gathers the amount of disk storage visible based on the computers being monitored by the Storage Resource agent or Tivoli Storage Productivity Center for Data agents. From the Monitored Computer Space report, you can view information about storage subsystems disks, known disks with serial numbers, or unknown disks without serial numbers.

Go through the following steps to get your Monitored Computer Space:

 In the Navigation Tree go to Data Manager → Reporting → Monitored Computer Storage Space → Disk Storage → All Disk Storage → By Computer, as shown in Figure 15-36.



Figure 15-36 By Computer

2. In the Selection panel, select **Distributed Computer Disk Space** option from the Relate Computers To drop-down menu. Select the columns shown in Figure 15-37.

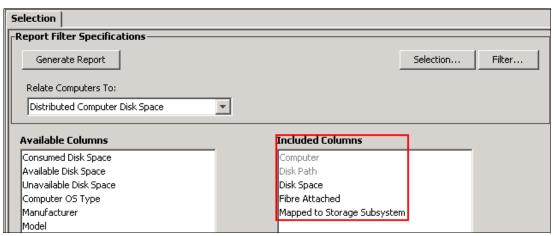


Figure 15-37 Data Licensing, columns needed for the report

 Click Generate Report. The output (Figure 15-38) shows that Tivoli Storage Productivity Center is managing 2.16 TB. This value is the amount of disk storage currently visible to computers with a Storage Resource agent or Tivoli Storage Productivity Center for Data agent installed.

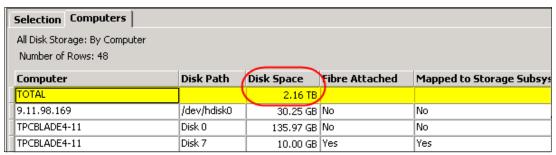


Figure 15-38 Data Licensing, generated report

15.5.3 Tivoli Storage Productivity Center Standard Edition licensing

Finally, we describe Tivoli Storage Productivity Center Standard Edition licensing, which is a combination of the two scenarios: Tivoli Storage Productivity Center for Data and Tivoli Storage Productivity Center for Disk. This edition is the total amount of Formatted Storage visible by Tivoli Storage Productivity Center for Disk and Monitored Computer Space visible by Tivoli Storage Productivity Center for Data, taking into account the storage already correlated or counted by Tivoli Storage Productivity Centerfor Disk such as IBM N Series.

Calculate Tivoli Storage Productivity Center Standard Edition monitored space:

- 1. Follow the steps shown for Tivoli Storage Productivity Center for Disk Licensing (see Figure 15-34 on page 590). In our example we got 103.33 TB.
- 2. Follow the steps shown for Tivoli Storage Productivity Center for Data Licensing but at the second step, on the "Included Columns" make sure you add "Correlated" as shown in Figure 15-39.

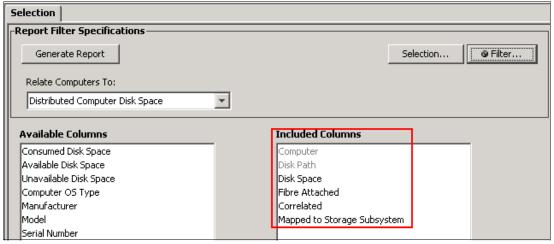


Figure 15-39 Tivoli Storage Productivity Center SE Licensing, columns needed for the report

3. On the same Selection panel, click **Selection** to deselect all NAS storage being managed through SMI-S. A sample panel is shown in Figure 15-40. We do this step because we do not want to double-count NAS Storage that is already being managed by Tivoli Storage Productivity Center for Disk. Click **OK** to close this panel.

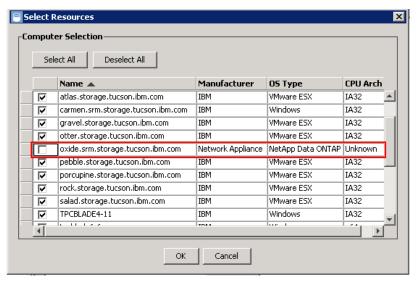


Figure 15-40 Tivoli Storage Productivity Center SE, Computer Selection panel

4. Click Filter button. The Edit Filter panel opens (Figure 15-41). From this panel, set Column to Correlated, set Operator to equals sign (=), and set Value 1 to No. This setting is required because using Yes would include the storage already counted by Tivoli Storage Productivity Center for Disk and we do not want to count this storage.

Therefore we select **Correlated = No**, which is data that Tivoli Storage Productivity Center for Data manages but Tivoli Storage Productivity Center for Disk does not. Click **OK** to close the panel and return to the panel shown in Figure 15-39 on page 592.

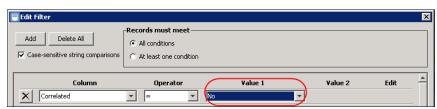


Figure 15-41 Tivoli Storage Productivity Center SE Licensing, Edit Filter panel

Click Generate Report from the Selection panel. The output is similar to Figure 15-42.
 Now, in the first row of the Disk Space column lists the amount of non-correlated disk storage visible to Tivoli Storage Productivity Center for Data. In this example it is 1.69 TB.

Selection Computers								
All Disk Storage: By Computer Number of Rows: 41								
Computer	Disk Path	Disk Space	Fibre Attached	Correlated	Mapped to 9			
TOTAL		1.69 TB	1					
9.11.98.169	/dev/hdisk0	30.25 GB	No	No	No			
TPCBLADE4-11	Disk 0	135.97 GB	No	No	No			
TPCBLADE4-11	Disk 7	10.00 GB	Yes	No	Yes			
TPCBLADE4-11	Disk 8	1.00 GB	Yes	No	Yes			

Figure 15-42 Tivoli Storage Productivity Center SE Licensing, generated report

Finally, the total amount of data is the sum of Tivoli Storage Productivity Center for Disk and Tivoli Storage Productivity Center for Data non-correlated. In our example 105.02 TB (Figure 15-43).

TPC for Disk	103.33 TB	Step 1
TPC for Data +	1.69 TB	Step 5
TPC Standard Edition =	105.02 TB	

Figure 15-43 Tivoli Storage Productivity Center Standard Edition managed capacity

15.6 VDisk mirroring reporting

In this section, we show you how to list which SVC or Storwize V7000 VDisks are mirrored. A basic example shows you how, from the GUI, you can see which volume defined in SVC or Storwize V7000 has a mirror copy and see all its attributes from the database views available. First, become familiar with Tivoli Storage Productivity Center v4.2 views at the following address:

http://www.ibm.com/support/docview.wss?uid=swg27019483

From Tivoli Storage Productivity Center GUI go to **Data Manager** \rightarrow **Reporting** \rightarrow **Asset** \rightarrow **By Storage Subsystem** and select your SVC or Storwize V7000 subsystem. In our examples we use SVC, but for Storwize V7000 the information is the same. You can see the available menu options as shown in Figure 15-44 on page 595. For a single mirrored volume, the VDisk and its mirror reside on separate mdisk groups. So, in Figure 15-45 on page 595, we show you the VDisk and its mirrored copy. Note that this VDisk has only one mirror. For this example, we show only the configuration for _32_mb_vd0409 VDisk. The GUI does not offer a direct way to get this information. From the database views, using several SQL statements you can get this information readily.

Warning: Insertion of new data or modification of existing data in the repository tables is currently out of scope and generally not possible. Attempting to do so can cause corruption to the data base repository. The only supported way of accessing the repository views is for read-only SELECT statements.

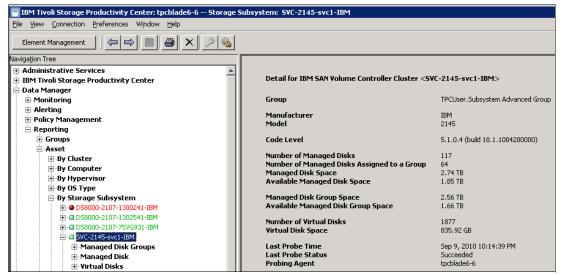


Figure 15-44 SVC Storage Subsystem menu

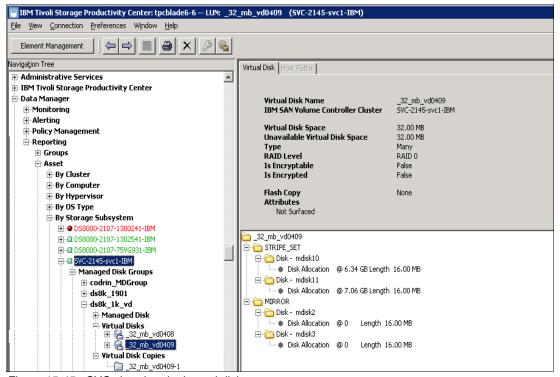


Figure 15-45 SVC virtual and mirrored disks

The steps are as follows:

1. Identify your target storage virtualization subsystem.

From the STORAGESUBSYSTEM view we are interested in the SUBSYSTEM_ID column which we use to filter the VDisks later. Part of this view is shown in Figure 15-46.

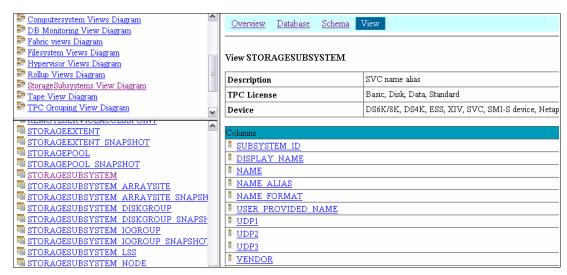


Figure 15-46 TPCDB Data Model, STORAGESUBSYSTEM view

The SQL statement and the associated output are shown in Figure 15-47.

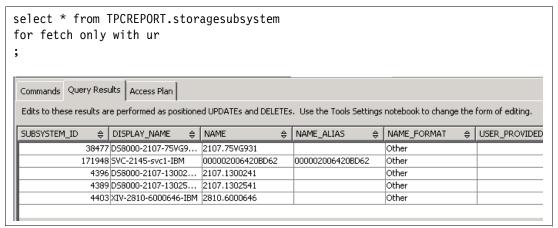


Figure 15-47 SQL statement and Storage Subsystem output for the Subsystems ID

Important: To avoid locking issues and avoid interfering with the Tivoli Storage Productivity Center server, always use an isolation level of UR (Uncommitted Read) and a "fetch only" connection type. Appending the following line to any SQL statement sets that information:

for fetch only with UR

2. Select the volumes from the Storage Subsystem.

After you have your Storage Subsystem ID, you might also be interested to see the VDisks that are mirrored. This query is based on the VDISK_COPY view, which shows the VDisks copy information. For VDisks mirror copies, VDISK_COPY_TYPE column must be equal to Secondary. See Figure 15-48.

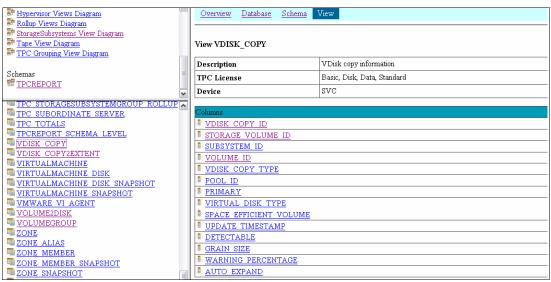


Figure 15-48 TPCDB Data Model, VDISK_COPY view

The SQL statement and its output are shown in Figure 15-49. The output shows you all the VDisks for the SVC-2145-svc-IBM, ID 171948, that are the mirror copy of some VDisk on that SVC. In the STORAGE_VOLUME_ID column, the ID will be equal to the VDisk source of the mirror relationship. In example, considering the first line of the output, you can verify this by querying the STORAGEVOLUME view filtering by ID=173268. This ID matches vdisk _32_mb_vd0409.

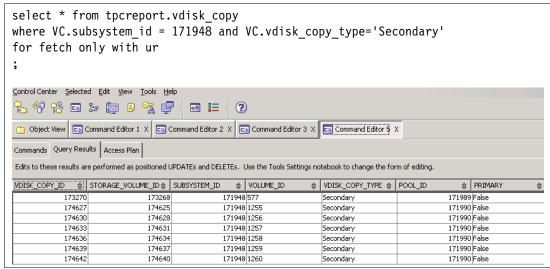


Figure 15-49 SQL statement and Storage Subsystem output for the virtual disks copy

3. Join several tables.

As a final step, we join several tables to show the VDisks that are mirrored. We filter and show only some of the columns including the STORAGE_VOLUME_ID, DISPLAY_NAME from the STORAGEVOLUME view and the DISPLAY_NAME from the STORAGEPOOL view, among others. The resulting SQL is shown in Figure 15-50.

```
select
   SV.storage_volume_id as VDISK_ID,
   SV.display_name as VDISK_NAME,
   SV.pool_id as VDISK_POOLID,
   SP.display_name as VDISK_POOLNAME,
   SPC.pool_id as VDISK_COPY_POOLID,
   SPC.display_name as VDISK_COPY_POOLNAME
        from tpcreport.storagevolume SV
full join tpcreport.vdisk_copy VC on SV.storage_volume_id=VC.storage_volume_id
full join tpcreport.storagepool SPC on VC.pool_id=SPC.pool_id
full join tpcreport.storagepool SP on SV.pool_id=SP.pool_id
where VC.subsystem_id = 171948 and VC.vdisk_copy_type='Secondary'
for fetch only with ur
;
```

Figure 15-50 SQL statement for joining different tables

See Figure 15-51. To explain further about the SQL statement, we joined the tables STORAGEPOOL, STORAGEVOLUME, and VDISK_COPY. From the STORAGEVOLUME, we get the information for the VDisks that are mirrored joining with the VDISK_COPY where the mirror copies reside. Finally, joining with the STORAGEPOOL gives you information about where the VDisks are defined. A final comment is that the intention of this section is to provide you with a starting point and, for example, only works for the VDisks with one mirror copy.

Control Center Selected Edit View Iools Help 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1								
Object View Command Editor 1 X Command Editor 2 X Command Editor 3 X Command Editor 5 X								
Commands Query Results Access Plan								
Edits to these results are performed as positioned UPDATEs and DELETEs. Use the Tools Settings notebook to change the form of editing.								
VDISK_ID	VDISK_NAME	VDISK_POOLID	VDISK_POOLNAME ♦	VDISK_COPY_POOLID ⇔	VDISK_COPY_POOLNAME ♦			
173268	_32_mb_vd0409	171990	ds8k_1k_vd	171989	ds8k_1901			
174625	vali_0	171991	eugen_test	171990	ds8k_1k_vd			
174628	vali_1	171991	eugen_test		ds8k_1k_vd			
174631	vali_2	171991	eugen_test		ds8k_1k_vd			
174634	vali_3	171991	eugen_test	171990	ds8k_1k_vd			

Figure 15-51 SQL output for the join statement

15.7 Storwize V7000 and SVC internal disk reports

In this section, we introduce the Tivoli Storage Productivity Center 4.2 enhancements for Storwize V7000 and SVC internal (local) disk reports. For our examples in this section, we focus on Storwize V7000 because the same concepts apply to SVC with internal SSDs. Support for these reports has been added in two sections:

- ▶ Data Manager → Reporting → Asset → By Storage Subsystem
- ▶ Disk Manager \rightarrow Reporting \rightarrow Storage Subsystems \rightarrow Disk \rightarrow By Storage Subsystem

In order to collect the data needed for these reports, you must perform the following actions:

- 1. Add a Storwize V7000 or SVC device with internal disks through the Tivoli Storage Productivity Center Configuration Wizard.
- 2. Run a successful subsystem probe job.
- 3. Generate the reports.

15.7.1 Data Manager report

Here is how you generate the Storwize V7000 internal disk report available within the Data Manager (see Figure 15-52).

- 1. Navigate to: Data Manager \rightarrow Reporting \rightarrow Asset \rightarrow By Storage Subsystem
- 2. Expand the node for your Storwize V7000 device.
- 3. Expand the Internal Disks node.
- 4. Left-click the disk to display the report with the disk details.

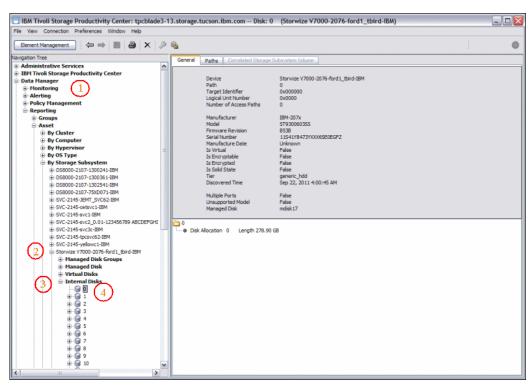


Figure 15-52 Steps to generate the Storwize V7000 internal disk report in Data Manager

15.7.2 Disk Manager report

Here we show you the Storwize V7000 internal disk report available within the Disk Manager:

- 1. Navigate to: Disk Manager \rightarrow Reporting \rightarrow Storage Subsystems \rightarrow Disk \rightarrow By Storage Subsystem
- 2. Click the Selection button.
- 3. Select the Storwize V7000 device(s) for which you want the report and click OK.

4. Click the **Generate Report** button as seen in Figure 15-53.

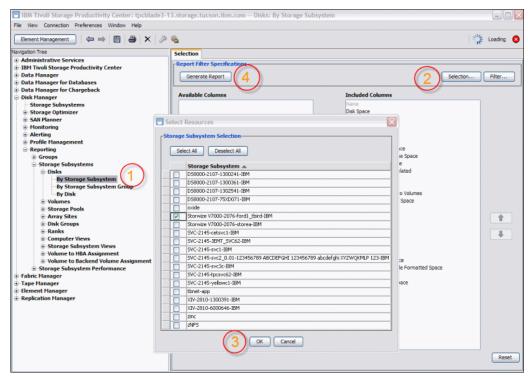


Figure 15-53 Steps to generate the Storwize V7000 internal disk report in Disk Manager

5. Click the magnifying glass icon to the left of the device to drill down in Figure 15-54.

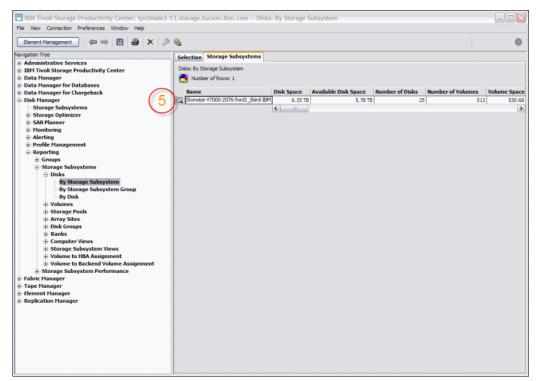


Figure 15-54 Click the magnifying glass icon next to the subsystem to list the internal disks

6. Click the magnifying glass icon to the left of the internal disk to get its detailed data as shown in Figure 15-55.

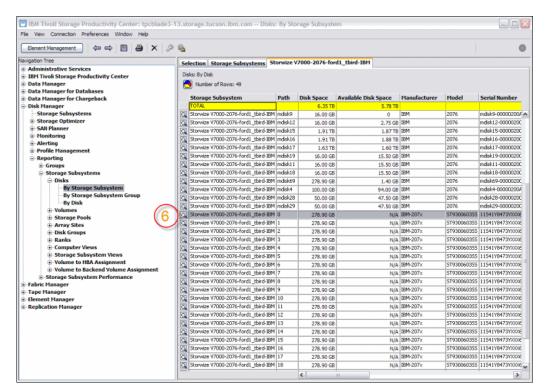
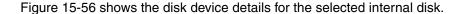


Figure 15-55 Click the magnifying glass icon next to the internal disk to get disk device details



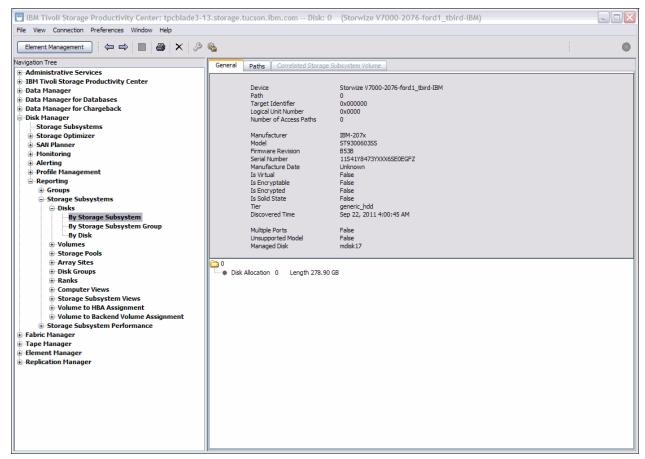


Figure 15-56 Internal disk device detail data for Storwize V7000

15.7.3 Troubleshooting Internal disk reports

If you encounter problems with these reports, the problem is most likely in one of two areas:

- ► Internal Disks is not shown under the Managed Disk node in Data Manager reports:
 - Internal disks will only appear under the Managed Disk node if the Mdisk is from internal disks.
- ► Internal Disks are not included in Disk Manager reports:
 - Check the probe job and make sure that it completed without any errors.
 - Check if there are any internal disks in the Storwize V7000 or SVC using the element manager GUI.

15.8 Storage Tier reports

This section shows you a new reporting capability introduced in Tivoli Storage Productivity Center V4.2.2 called Tivoli Storage Productivity Center Tier Reports, referenced from now on as *Storage Tier reports*.

15.8.1 Tivoli Common Reporting and Storage Tier reports

The 4.2.2 release of Tivoli Storage Productivity Center introduces Tivoli Common Reporting 2.1.1, which uses the IBM Cognos® reporting engine. This is the foundation for future planned enhancements to reporting in Tivoli Storage Productivity Center.

Tivoli Common Reporting (TCR) is an additional component that must be installed separately in addition to Tivoli Storage Productivity Center. It can be installed on the Tivoli Storage Productivity Center server, or for large environments with a busy Tivoli Storage Productivity Center server we suggest a remote install on a separate server. For more information about the TCR installation on a separate server, refer to Appendix E, "Tivoli Common Reporting: Remote Installation" on page 841.

This release of TCR adds support for Storage Tier reports, which are described in detail in 15.8.4, "Running Storage Tier reports" on page 621 and in 15.8.5, "Analyzing Storage Tier reports: Case Studies" on page 633. Storage Tier reports can help you answer questions and provide new function such as these:

- ► Where should a new application be placed?
- When should a workload migration be considered between storage pools?
- ▶ Which workload should be migrated or re-tiered out of a storage pool?
- Which storage pool should be considered as target for a workload subject to be migrated?
- Storage Resource Group profiling
- ► Speed up "time to service" for new workloads

Tivoli Common Reporting uses a connection to the Tivoli Storage Productivity Center TPCDB database to extract the performance data that has been collected to generate these reports. This connection is depicted in Figure 15-57.

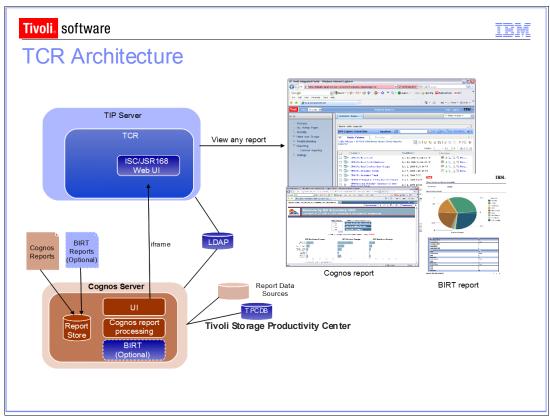


Figure 15-57 TCR architecture and integration with Tivoli Storage Productivity Center

15.8.2 Terminology

In this section we define the following terms.

Storage virtualizers

In this book, the term storage virtualizer is used to refer to SAN Volume Controller and Storwize V7000 storage systems.

Tiering

Tiering is the term that is used to describe storage systems that comprise different performing subsystems, or physical disks within a subsystem that have different price points. The subsystems and physical disks are organized into different MDisk groups (also referred to as pools) in SVC and Storwize V7000.

Up-tiering and down-tiering

Up-tiering and *down-tiering* are the terms that are used to refer to the process of moving a virtual disk (VDisk, also referred to as volume) from one type of MDisk group to another MDisk group. Up-tiering refers to the movement from a lower capability pool to a higher capability pool. Usually, this means going from a lower cost pool to a higher cost pool, but that is not always the case.

Down-tiering refers to the movement of a workload from a higher capability pool to a lower capability pool. Usually this means going from a higher cost pool to a lower cost pool, but again, this is not always the case.

Capability versus capacity

Capability is the term that is used to refer to performance, whereas capacity is the term that is used to refer to space. Storage tier reports help users to engage in capability planning to ensure that existing and future workloads meet or exceed the service level agreements that govern the interaction between storage and applications. Traditionally, users have focused on capacity planning to ensure that they have sufficient space for existing and future applications.

15.8.3 TCR installation and setup

This section explains the installation of Tivoli Common Reporting (TCR) on Windows 2008. We depict the steps needed to deploy the software and show you how to configure TCR to connect to the Tivoli Storage Productivity Center database on the same system.

TCR can also be installed on AIX and Linux. The steps are the same, except that instead of running '.exe' or '.bat' command files, you will be running '.sh' scripts that have the same name and use the same options (with adjustments for AIX/Linux path names).

Before you start

Make sure your server meets the system requirements for TCR. If you are installing TCR on your Tivoli Storage Productivity Center server, you need to make sure that the minimum system requirements for all software components *considered together* are met. For example, if component A requires a minimum of 4 GB and component B requires a minimum of 4 GB, then your system should have a minimum of 8 GB if both A and B are to be installed on the same server. For system requirement details, see 2.2.1, "Verifying system hardware and software prerequisites" on page 25.

To successfully complete the TCR deployment you need the file TPC_Reports.zip file from the \tool\report subdirectory of your Tivoli Storage Productivity Center installation source. Make sure it is available.

For TCR, the required version of the IBM Java runtime environment is included in deployment package. You can also find the installation files for the IBM Java runtime environment if you open a web browser to the Tivoli Storage Productivity Center server at this website:

http://tpcserver:9550/ITSRM/app/welcome.html

Installing TCR on your server

Follow these steps:

 To start the installation, run the launchpad.exe executable as shown in Figure 15-58 to start the IBM Tivoli Common Reporting Installation assistant. On AIX and Linux, run the command ./launchpad.sh.

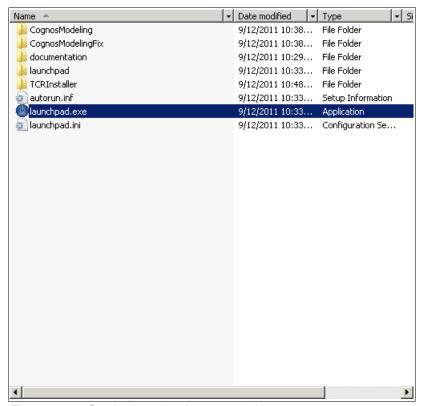


Figure 15-58 Starting launchpad.exe executable

Important: Make sure to have a supported web browser installed in your system *before* starting the TCR Installation assistant. See the "Hardware and Software Requirements" section in the TCR Information Center for supported web browsers, as well as supported hardware and software requirements, at the following website:

http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr .doc 211/rtcr soft and hard regs.html

2. The TCR Installation assistant prompts you with an overview of available options. Select **Install IBM Tivoli Common Reporting** as shown in Figure 15-59.

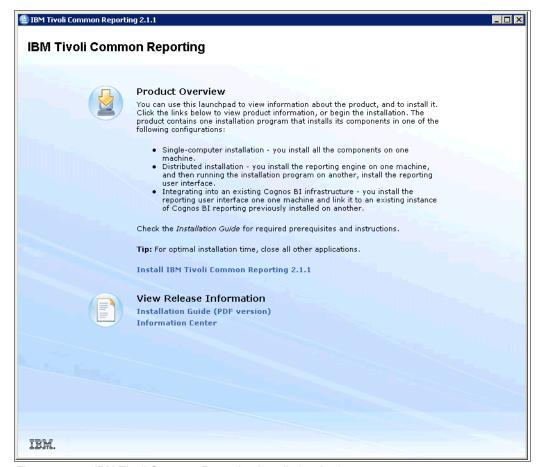


Figure 15-59 IBM Tivoli Common Reporting Installation Assistant

3. The installation will begin and deploy setup files to the system as shown in Figure 15-60.

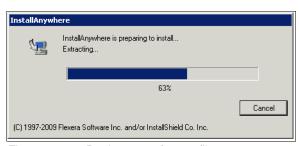


Figure 15-60 Deployment of setup files

4. Choose a language from the drop-down menu of the language selection dialog of the TCR installation as shown in Figure 15-61. Click **OK** to continue with the installation.

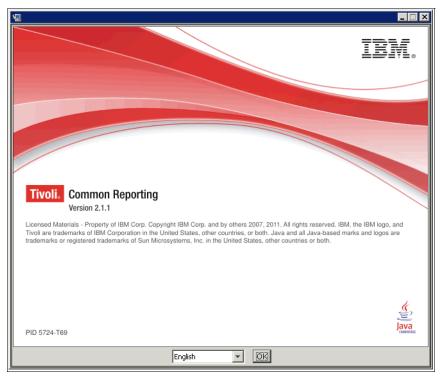


Figure 15-61 Language selection dialog

5. The TCR installation assistant welcomes you with a welcome message as shown in Figure 15-62. Click **Next** to continue.

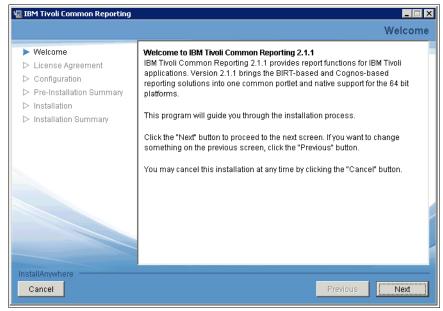


Figure 15-62 TCR installation assistant welcome message

6. Decide on whether or not you can accept the license agreement and enter your choice in the software license agreement dialog shown in Figure 15-63. Accepting the software license agreement will allow you to click **Next** to continue.

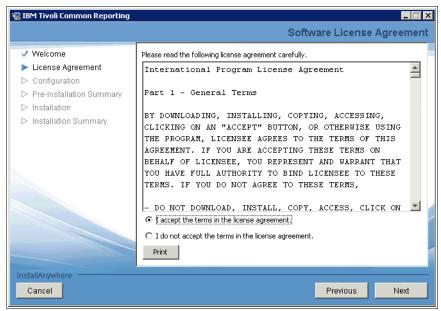


Figure 15-63 TCR installation assistant software license agreement dialog

7. The TCR installation assistant will initialize the Deployment Engine (DE) as shown in Figure 15-64. Wait for it to complete.

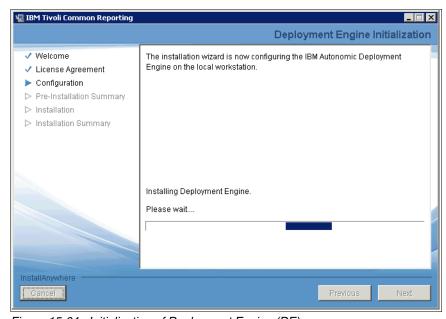


Figure 15-64 Initialization of Deployment Engine (DE)

8. Select the "Install a new instance of Tivoli Common Reporting" of the installation modes presented to you in Figure 15-65. Afterwards click **Next** to continue.

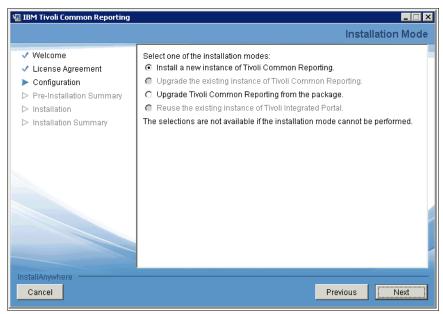


Figure 15-65 Installation Mode selection panel

 Select the appropriate installation scenario from the Installation Scenario Selection shown in Figure 15-66. In this example we will show the "Single-computer installation" variant of the possible TCR deployment strategies. Click **Next** to continue after making your selection.

Tip: For large scale deployments, consider installing the TCR reporting component onto a machine different from your Tivoli Storage Productivity Center server. This will allow you to distribute the load across multiple servers which will result in higher overall performance and responsiveness. For more information about the TCR installation on a separate server, refer to Appendix E, "Tivoli Common Reporting: Remote Installation" on page 841

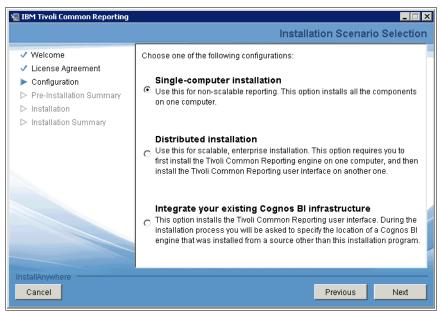


Figure 15-66 Installation Scenario Selection

10. Select your installation directory in the "Installation Directory Selection" panel shown in Figure 15-67. We suggest the default of C:\IBM\tivoli\tcr (/opt/IBM/tivoli/tcr on AIX and Linux). Do not use spaces or special characters in the path name. Click **Next** to continue.

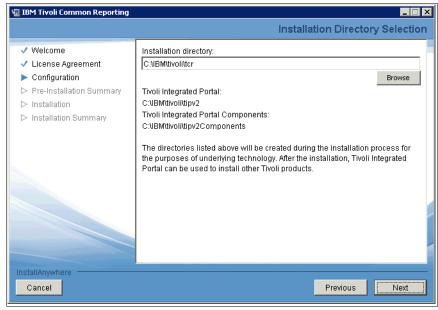


Figure 15-67 Installation Directory Selection

11. Provide credentials for your WebSphere infrastructure as shown in the "WebSphere Information" dialog seen in Figure 15-68. The Port Number you select here will be used as the initial port for a port range of 15 ports. If you are installing TCR on the Tivoli Storage Productivity Center server with an existing Tivoli Information Portal (TIP) instance, port 16310 will already be used by the first TIP instance. If this is the case, we suggest that you use 17310 for the TCR TIP instance. Click **Next** to continue.

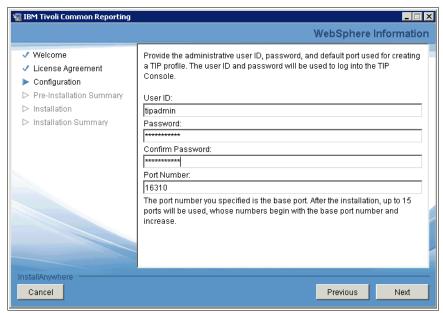


Figure 15-68 WebSphere Information

12.If the port you choose is in use, the installer will let you reconsider your choice or allow you to continue with the automatically created proposal as shown in Figure 15-69. Click **Next** to continue.

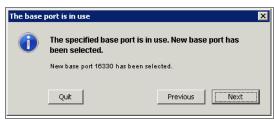


Figure 15-69 The base port is in use

13. Provide a port number for the IBM Cognos Content Database as shown in Figure 15-70. Choose your own port or stick with the default port which is 1527. Click **Next** to continue.

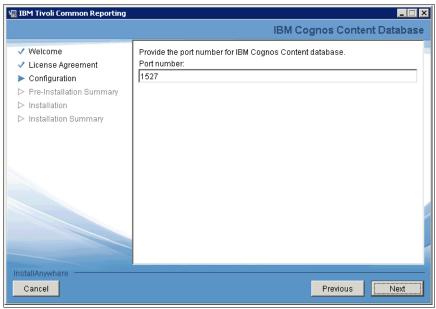


Figure 15-70 IBM Cognos Content Database

14. The TCR installation assistant will now prepare the installation as shown in Figure 15-71. Wait for it to be completed.

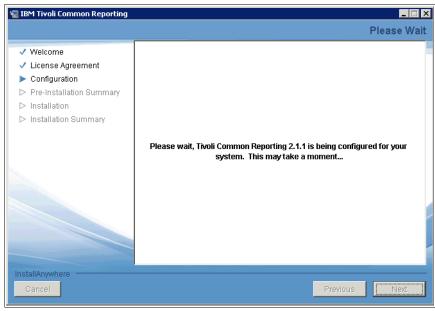


Figure 15-71 Installation preparation takes place

15. After successful installation preparation the TCR installation assistant will display an overview of the choices you just made. Review your Pre-Installation Summary as shown in Figure 15-72. If you need to make a change, click 'Previous' and return to the respective panel to change the values. If you agree with the provided overview click **Install** to start the installation.

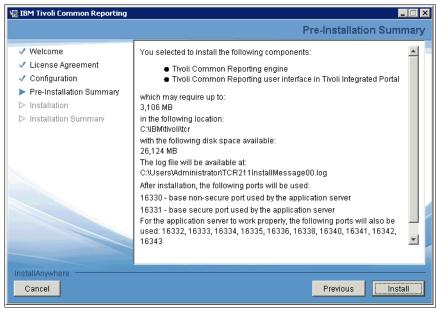


Figure 15-72 Pre-Installation Summary

16. The installation will begin. While you monitor the progress bar shown in Figure 15-73 you can view the Windows Resource Monitor during the installation. Should the Memory section of the Windows Resource Monitor contain a number above zero for the Hard Faults/min column, this could be an indication for your system being memory constrained.

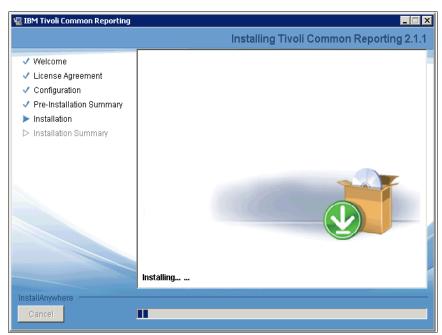


Figure 15-73 TCR Installation running

17. After the install completes, you will see the Installation Summary as shown in Figure 15-74 on page 616. Review the information provided and take note of the URL needed to start using your TCR instance. It will look something like this:

http://tpcblade3-13.storage.tucson.ibm.com:16330/ibm/console

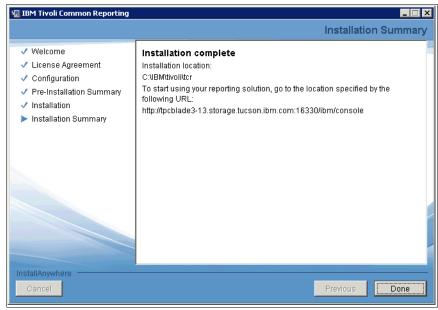


Figure 15-74 Installation Summary

You need to use a supported web browser on the machine where you will be launching TCR (this might be a machine other than the server where you installed TCR). Refer to the 'Supported Platforms' document for Tivoli Storage Productivity Center at:

https://www-304.ibm.com/support/docview.wss?uid=swg27019380#browser

You can also refer to the 'Hardware and Software Requirements' section in the TCR Information Center for supported web browsers at:

http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/rtcr_soft_and_hard_reqs.html

Before you can create reports you need to connect the newly installed TCR to the Tivoli Storage Productivity Center database.

18.To connect the newly installed TCR with the Tivoli Storage Productivity Center database you need to create a data source connection using the TCR command line tool. Open up a windows command line window by clicking **Start** → **Run**, enter **cmd** as shown in Figure 15-75 and click **OK** to continue. On AIX and Linux you can run the command from a shell window (xterm, etc.).



Figure 15-75 Opening a windows command line window

19. Change your current working directory to the location of the TCR command line tool. On AIX and Linux, the path is under /opt/IBM/tivoli/... if you used the suggested default, and on Windows the command to change to the default location is:

cd c:\IBM\tivoli\tipv2Components\TCRComponent\bin

Tip: The following commands are long and need to be typed in manually. Double-check your input on these commands. A simple typo, the accidental use of an equals sign, or a misplaced space will lead to a hanging command that might not return with a message. You might need to use CTRL+C to abort the command and retry.

20. The first of the three steps to connect TCR to your Tivoli Storage Productivity Center database is to create a data source connection. Enter the following command with the necessary changes for your environment.

trcmd.bat -user tipadmin -password tippassword -dataSource -add TPCDB -dbType DB2 -connectionName TPCDB -dbName TPCDB -dbLogin db2admin -dbPassword db2password -groups Everyone

Attention: AIX and Linux uses ./trcmd.sh with the same parameter string.

The successful completion of the above command is indicated with the following message: CTGTRQ113I Data source TPCDB was successfully added/modified.

21. After creating your data source connection, we suggest you test the data source connection. Use the following command to verify your data source connection:

trcmd.bat -user tipadmin -password tippassword -dataSource -test TPCDB -dbLogin db2admin -dbPassword db2password

The successful completion of the data source connection test is indicated with the following message:

CTGTRQ114I Data source TPCDB was successfully tested.

Tip: If the test of your data source connection fails, ensure that your DB2 instance is running.

22. Finally we import the deployment package into TCR. The deployment package is a *.zip file that contains the Tivoli Storage Productivity Center reports. The file is named TPC_Reports.zip. It is provided with the Tivoli Storage Productivity Center installation source and resides in the tool\reports subdirectory of it. Run the following command to import the deployment package:

trcmd.bat -import -bulk c:\TPC4.2.2.74-windows\tool\report\TPC_Reports.zip -user tipadmin -password tippassword

Attention: Avoid spaces and special characters in the path name to your deployment package. Otherwise the import will fail.

The successful completion of the deployment package import is indicated by the following message:

CTGTRQ092I Import operation successfully performed.

23. Configuration is complete. You can now log on to your TCR instance by pointing your web browser to the URL noted earlier:

http://tpcblade3-13.storage.tucson.ibm.com:16330/ibm/console/

You might be prompted with a message similar to this one: "There is a problem with this website's security certificate." See Figure 20-4 on page 788 for details. Choosing to ignore the exception and continue (the exact steps will vary depending on the web browser you are using) will allow you to continue. To logon to TCR as shown in Figure 15-76, you need to enter the credentials you provided earlier during the installation. Enter the User ID and Password and click **Log in** to continue.



Figure 15-76 TCR Log in panel

24.After successfully logging on to your TCR instance, you need to click **Reporting** on the left-hand side of the panel shown in Figure 15-77, which shows the navigation bar. This will allow you to access the Common Reporting option.



Figure 15-77 TCR Welcome window

25. Click **Common Reporting** on the left hand side of the panel as shown in Figure 15-78 to access the actual Tivoli Storage Productivity Center reporting.



Figure 15-78 TCR Welcome window with Common Reporting expanded

26.Click **TPC** to access the functionality of the IBM Tivoli Common Reporting as shown in Figure 15-79.

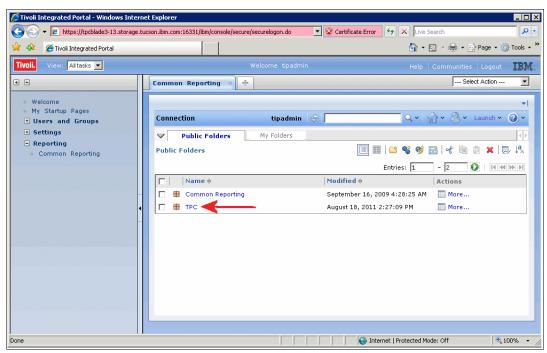


Figure 15-79 TCR TPC integration

27. From the panel shown in Figure 15-80 you can now access the Tivoli Storage Productivity Center reports available with TCR. We suggest that you click the small black arrow pointing to the left (◀) on the bar that separates the left hand side navigation menu from the right hand side report overview. By doing this you will maximize the display space for your reports.

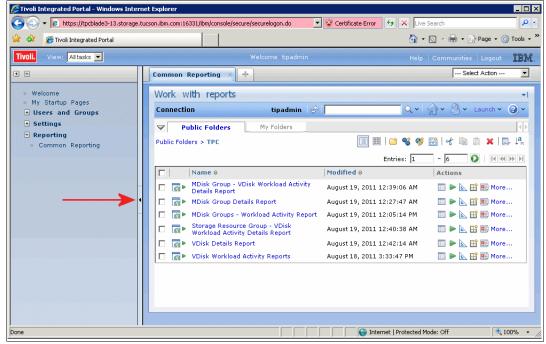


Figure 15-80 TCR Reports List

28. The resulting panel will look like Figure 15-81. You are now ready to create reports.

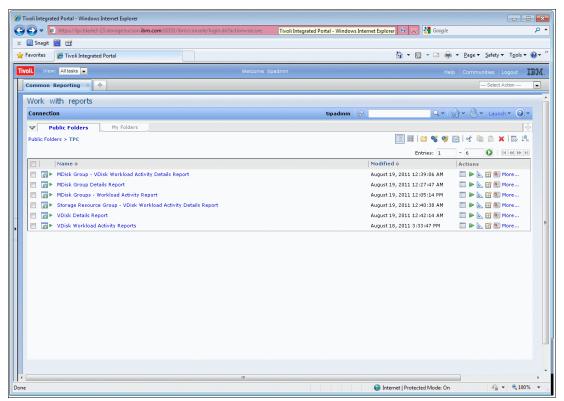


Figure 15-81 TCR Reports List without left hand side navigation menu

Tip: Right-click the page and Refresh or press F5 to refresh the screen and fix any display challenges you might face. Remember: *Do not* use the browser's **back** button.

15.8.4 Running Storage Tier reports

With all the Storage Tier reports being available from the TCR Reports List you are now ready to run the reports. Feel free to work with Storage Tier reports as you see fit or follow any of our examples below. The list of predefined reports consists of the following:

- ► MDisk Group VDisk Workload Activity Details Report
 - Top 5 VDisks from a selected MDisk Group
 - Use the MDisk Group VDisk Workload Activity Details Report to monitor the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity of VDisks in an MDisk Group is determined by calculating the average peak utilization of each VDisk.
- Mdisk Group Details Report
 - Accumulated statistics for selected MDisk Group
 - Use the MDisk Group Details Report to monitor the workload activity of a managed disk (Mdisk) group
- MDisk Groups Workload Activity Report
 - Accumulated statistics for all MDisk Groups of selected Virtualizer

- Use the MDisk Groups Workload Activity Report to monitor the workload activity of managed disk (Mdisk) groups
- ► Storage Resource Group VDisk Workload Activity Details Report
 - Top 5 VDisks from a selected SRG
 - Use the Storage Resource Group VDisk Workload Activity Details Report to monitor the workload activity of virtual disks (VDisks) in a Storage Resource Group. The workload activity of VDisks in a Storage Resource Group, such as a database application, is determined by calculating the average peak utilization of each VDisk.
- VDisk Details Report
 - Detailed statistics for selected VDisk
 - Use the VDisk Details report to monitor the workload activity of a virtual disk (VDisk) in a managed disk group
- VDisk Workload Activity Reports
 - Drill through capable reports containing statistics based on selection of MDisk,
 Virtualizer, or SRG

Working with Storage Tier reports from the back-end perspective

MDisk Group based reports are helpful to report on the back-end status of your environment.

 Starting from the Report Selection panel, as seen in Figure 15-81 on page 621, click MDisk Groups - Workload Activity Report. Select the Storage Virtualizer for which you want to run the report as shown in Figure 15-82.

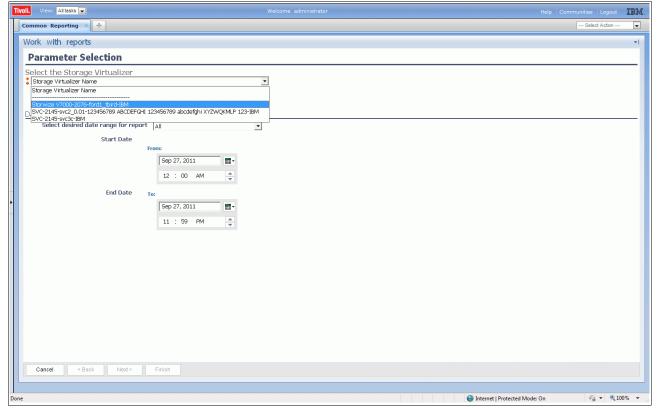


Figure 15-82 Storage Virtualizer Selection

2. Select the Date Range for your Report from the drop-down menu shown in Figure 15-83. Finally click **Finish** on the bottom of the page to start the report generation.

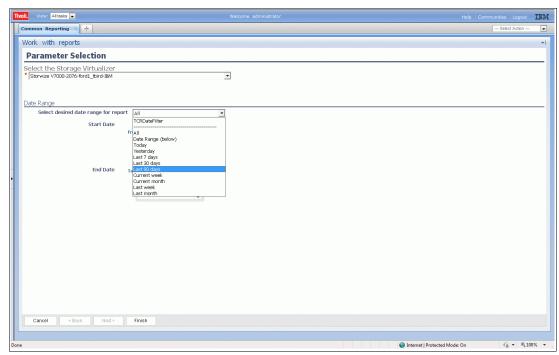


Figure 15-83 Date Range Selection

Tip: Be aware that the *Start Date* and *End Date* input fields are only used if you select **Date Range** from the drop-down menu. Otherwise they are without function. For example, if you select **Last 7 days**, you do not need to fill out these fields.

3. Your MDisk Groups - Workload activity Report will be generated and presented to you as shown in Figure 15-84. From here you can continue to click one of the bars from the bar chart or on one of the MDisk groups (Figure 15-84).

Clicking on a bar that represents an MDisk group will present the MDisk Group - VDisk Workload Activity Details Report (Figure 15-92 on page 630), which identifies the most active and least active VDisks associated with that MDisk Group.



Figure 15-84 MDisk Groups - Workload activity Report

MDisk Groups - Details Report

Follow these steps

 Starting from the Report Selection panel, as seen in Figure 15-81 on page 621, click MDisk Groups - Details Report. Select the Storage Virtualizer for which you want to run the report as shown in Figure 15-85.

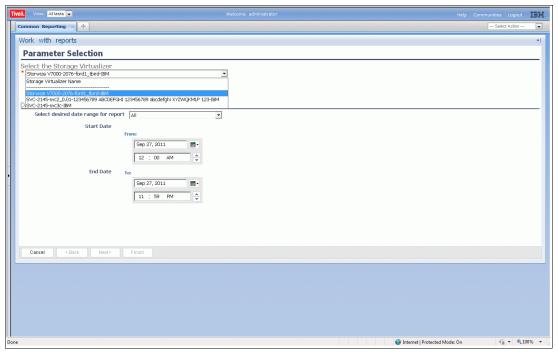


Figure 15-85 Storage Virtualizer Selection

2. Select the MDisk Group you want to generate the report for from the drop-down menu as shown in Figure 15-86.

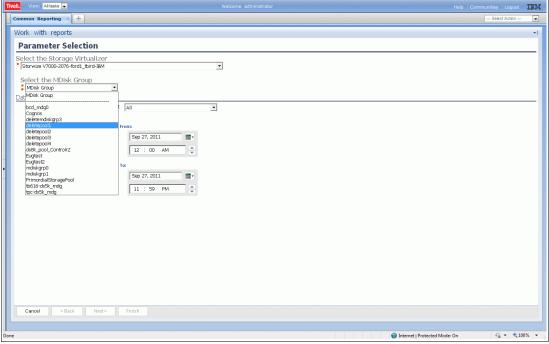


Figure 15-86 MDisk Group Selection

3. Select the date range for the report to be created from the drop-down menu as shown in Figure 15-87. Then click **Finish** to start the creation of your report.

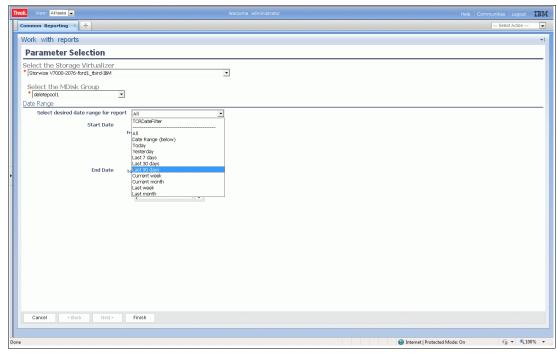


Figure 15-87 Date Range Selection

4. You will be presented the MDisk Group Details Report as shown in Figure 15-88.

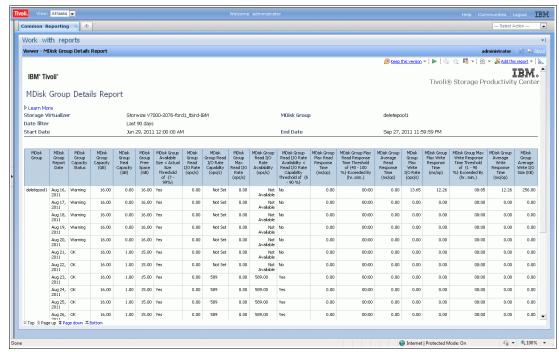


Figure 15-88 MDisk Group Details Report

Working with Storage Tier reports from a front-end perspective

VDisk based reports are helpful to report on the front-end status of your environment. In this example, we generate a report after using the following drill through navigation path:

- 1. MDisk Groups Workload Activity Report
- 2. MDisk Group VDisk Workload Activity Report
- 3. VDisk Details Report

Using the drill through navigation approach, you can easily follow the reports through your environment to identify potential improvement areas of your environment:

1. Starting from the Report Selection panel as seen in Figure 15-81 on page 621, start creating an **MDisk Groups - Workload Activity Report**. Select the Storage Virtualizer for which you want to run the report as shown in Figure 15-89.



Figure 15-89 Storage Virtualizer Selection

2. Select a date range for the report as shown in Figure 15-90.

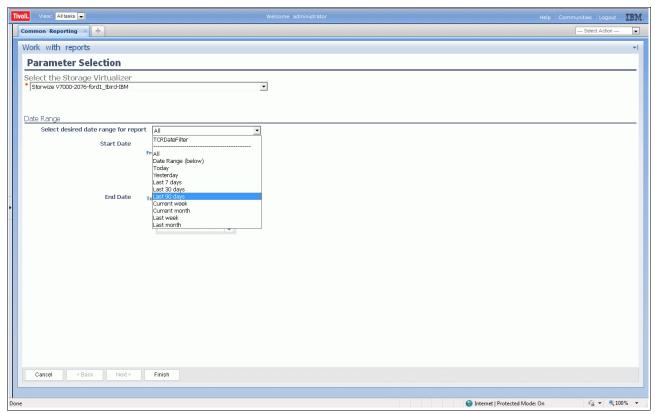


Figure 15-90 Date range selection

3. Your report will be generated and presented to you. From the MDisk Groups - Workload Activity Report, you can now identify the MDisk Group you are interested in and click the bar from the bar chart or the link as shown in Figure 15-91.



Figure 15-91 MDisk Groups - Workload Activity Report

4. The drill through navigation will present you with the MDisk Group - VDisk Workload Activity Details Report from the selected MDisk Group as shown in Figure 15-92. In the MDisk Group - VDisk Workload Activity Details Report, you can now easily spot your most active VDisks that are potential candidates for a migration to a higher performance storage tier. Also, at the same time you can identify your least active VDisks that potentially could be moved to a lower performance storage tier. Click any VDisk to use the drill through navigation to get more detailed information about it.

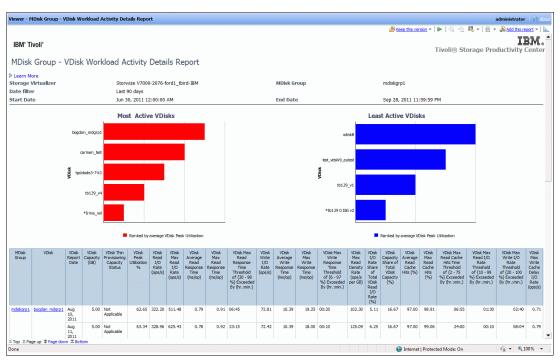


Figure 15-92 MDisk Group - VDisk Workload Activity Details Report

5. The VDisk Details Report of your selected VDisk will be presented to you to get more detailed information about it, as shown in Figure 15-93.

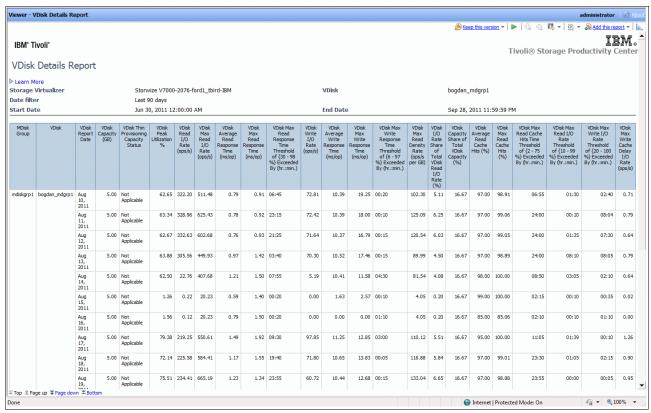


Figure 15-93 VDisk Details Report

Running a Storage Resource Group report

In IBM Tivoli Storage Productivity Center Version 4.2.2, the Storage Resource Group functionality has been extended to include reporting.

To run a Storage Resource Group report, follow these steps:

 When VDisk Workload Activity Reports is selected (Figure 15-81 on page 621), you are presented with three options. You can run a report that shows VDisks in MDisk Groups, VDisks in a storage virtualizer, or VDisks in Storage Resource Groups; Select Storage Resource Group - VDisk Activity Report and click Next (Figure 15-94).

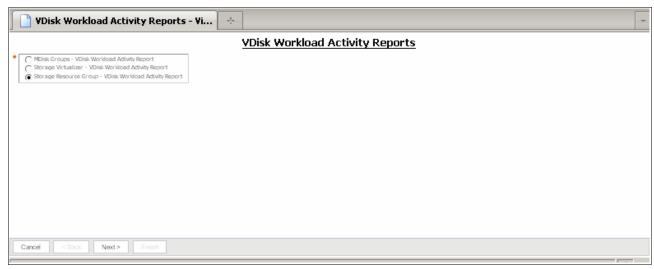


Figure 15-94 Selecting a Storage Resource Group - VDisk Workload Activity Report

2. Select a storage virtualizer and enter a date range or select a date filter as shown in Figure 15-95, then click **Finish**.

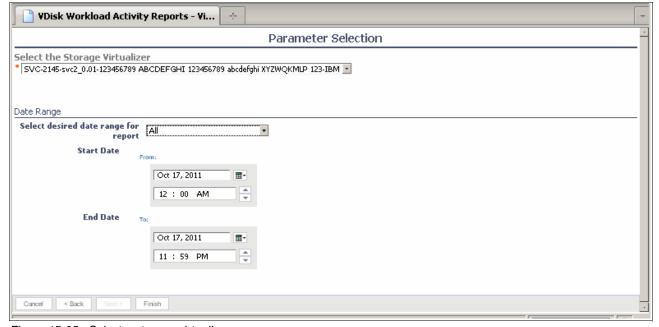


Figure 15-95 Select a storage virtualizer

3. A comparison is displayed that ranks the most active and least active Storage Resource Groups based on the average VDisk utilization (Figure 15-96).

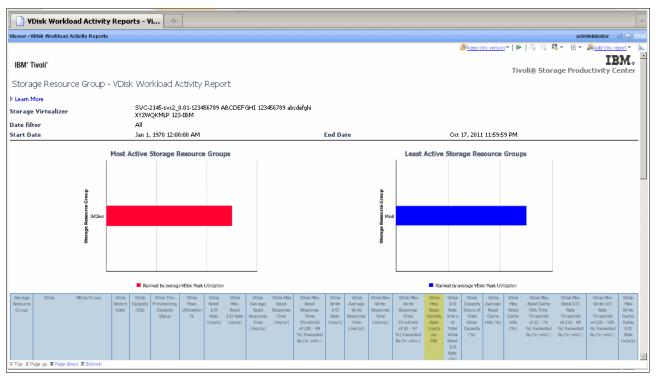


Figure 15-96 Storage resource groups in a storage virtualizer

15.8.5 Analyzing Storage Tier reports: Case Studies¹

In this section we provide two case studies that walk you through the process of analyzing the *MDisk Groups - Workload Activity Reports* to determine what Mdisk Groups are candidates for re-balancing ("Case Study 1: Identifying MDisk Groups for up-tiering and down-tiering") and the *VDisk Workload Activity Details Reports* to determine what VDisks within the managed disk groups are candidates for relocation ("Case Study 2: Identifying most active VDisks to reduce MDisk Group workload" on page 636).

Case Study 1: Identifying MDisk Groups for up-tiering and down-tiering

This case study describes how clients analyze the storage workloads of MDisk Groups to determine whether they can be migrated to other MDisk Groups in a storage virtualizer.

Overview

The starting point of the analysis procedure is the *MDisk Groups - Workload Activity Report* (Figure 15-97), which provides charts for the most active (and least active) MDisk Groups. The specific metrics for capacity and capability are then evaluated to determine whether further investigation is required like, for example, the response times. If it is determined that an MDisk Group requires re-balancing, then the VDisk report for the MDisk Group is opened (See "Case Study 2: Identifying most active VDisks to reduce MDisk Group workload" on page 636) where VDisk candidates for relocation are identified.

http://www-01.ibm.com/support/docview.wss?uid=swg27023263

¹ Materials and case studies for this section were taken directly from the IBM white paper, IBM Tivoli Storage Productivity Center Version 4.2.2 Storage Tier reports, at this website:

Steps for analyzing MDisk Groups

1. Use the most active or least active MDisk Group chart to select an MDisk Group that might require up-tiering or down-tiering (Figure 15-97).

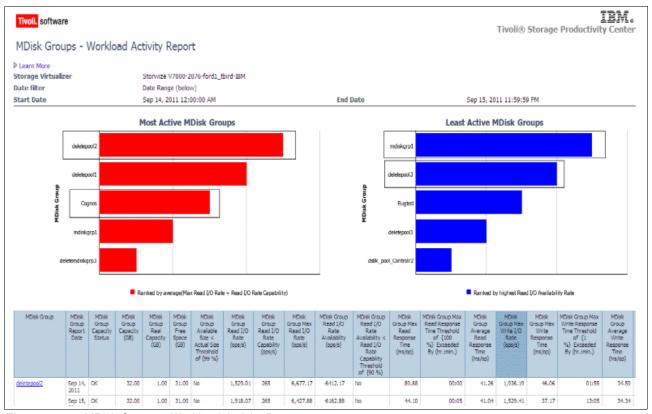


Figure 15-97 MDisk Groups - Workload Activity Report

 Analyze the values in the MDisk Group details chart for the most active or least active MDisk Group to determine if it is a candidate for up-tiering or down-tiering (Figure 15-98). In our example, the MDisk Group **deletepool2** is analyzed further to determine whether it is a suitable candidate for up-tiering.

MDisk Group Step 2- Analyze	sk Group MDisk Group Report Date	oup Group Gr port Capacity Cap ate Status (0		MDisk Group Real Capacity (GB)	MDisk Group Free Space (GB)	MDisk Group Available Size < Actual Size Threshold of {99 %}	MDisk Group Read I/O Rate (ops/s)	MDisk Group Read I/O Rate Capability (ops/s)	MDisk Group Max Read I/O Rate (ops/s)	MDisk Group Read I/O Rate Availability (ops/s)	MDisk Group Read I/O Rate Availability < Read I/O Rate Capability	MDisk Group Max Read Response Time (ms/op)	MDisk Group Max Read Response Time Threshold of {100 %} Exceeded By (hr.:min.)	MDisk Group Average Read Response Time (ms/op)	MDisk Group Max Write I/O Rate (ops/s)	MDisk Group Max Write Response Time (ms/op)	MDisk Group Max Write Response Time Threshold of {1 %} Exceeded By (hr.:min.)	MDisk Group Average Write Response Time (ms/op)
											Threshold of {90 %}	Step	3-Analyze the Re	ad/Write re	sponse tim	e and nos	of hours	
deletepool2	Sep 14, 2011	ОК	32.00	1.00	31.00	No	1,529.01	265	6,677.17	-6412.17	No	89.88	01:55	41.26	1,936.19	46.06	01:55	34.59
	Sep 15, 2011	ОК	32.00	1.00	31.00	No	1,918.07	265	6,427.88	-6162.88	No	44.10	13:05	41.04	1,529.41	37.17	13:05	34.34
deletepool2							1,723.54		Max 6,677.17	Avg -6,287.52		Max 89.88		Avg 41.15	Max 1,936.19	Max 46.06		Avg 34.47
deletepool 1	Sep 14, 2011	ОК	16.00	Step 1-/	Analyze	the Max Re	ad (50.82	589	6,420.02	-5831.02	No	69,40	12:00	39.24	1,585.84	42.87	12:00	30.37
	Sep 15, 2011	ОК	16.00	&availab		papility	0.00	589	0.04	588.96	Yes	81.80	00:20	63.39	0.00	0.00	00:20	0.00
deletepool1							Avg 575.41		Max 6,420.02	Avg -2,621.03		Max 81.80		Avg 51.32	Max 1,585.84	Max 42.87		Avg 15.19

Figure 15-98 Analyze MDisk Group deletepool2 for rebalancing

For the steps used, see Table 15-9, where the Mdisk Group **deletepool2** is analyzed further to determine whether it is a suitable candidate for up-tiering.

Table 15-9 Steps to determine if the deletepool2 MDisk Group is a candidate for up-tiering

Step	Parameters used to determine tiering	Observation	Up-tiering suggested
1	MDisk Group Read I/O Rate Capability, MDisk Group Max Read I/O Rate, and MDisk Group Read I/O Rate Availability	The maximum read I/O rate is high compared with capability, and read I/O rate availability is negative.	Yes
2	MDisk Group Capacity Status, MDisk Group Capacity, MDisk Group Real Capacity, and MDisk Group Free Space	The capacity status is OK. In fact, very little capacity is being consumed. Capability, on the other hand, is extremely over-utilized (the actual I/O rate maximum is over 25 times that of the theoretical capability).	Yes
3	MDisk Group Max Read Response Time, MDisk Group Max Read Response Time Threshold, MDisk Group Average Read Response Time, MDisk Group Max Write Response Time, and MDisk Group Max Write Response Time Threshold	Maximum and average read response times are high. Maximum and average write response times are high. The number of hours that the write response time exceeds the threshold is too high.	Yes
Final re	ecommendation	The MDisk Group <i>deletepool2</i> is over-utilized. Some of the workloads in this MDisk Group a candidates for up-tiering.	

Now we will also use the Most active MDisk Groups chart (Figure 15-97 on page 634) to determine if MDisk group **Cognos** is a suitable candidate for up-tiering. In Figure 15-99 the MDisk group **Cognos** is analyzed further to determine whether it is a suitable candidate for up-tiering.

MDisk Group	MDisk Group Report Date	MDisk Group Capacity Status	MDisk Group Capacity (GB)	MDisk Group Real Capacity (GB)	MDisk Group Free Space (GB)	MDisk Group Available Size < Actual Size Threshold of {99 %}	MDisk Group Read I/O Rate (ops/s)	MDisk Group Read I/O Rate Capability (ops/s)	MDisk Group Max Read I/O Rate (ops/s)	MDisk Group Read I/O Rate Availability (ops/s)	MDisk Group Read I/O Rate Availability < Read I/O Rate Capability Threshold of {90 %}	MDisk Group Max Read Response Time (ms/op)	MDisk Group Max Read Response Time Threshold of {100 %} Exceeded By (hr.:min.)	MDisk Group Average Read Response Time (ms/op)	MDisk Group Max Write I/O Rate (ops/s)	MDisk Group Max Write Response Time (ms/op)	MDisk Group Max Write Response Time Threshold of {1 %} Exceeded By (hr.:min.)	MDisk Group Average Write Response Time (ms/op)
deletepool2	Sep 14,	OK	32.00	1.00	31.00	No	1,529.01	265	6,677.17	-6412.17	OT {90 %}	89.88	01:55	41.26	1,936.19	46.06	01:55	34.59
	2011	10000	10,000,000	-							1015			-7			1	
	Sep 15, 2011	ОК	32.00	1.00	31.00	No	1,918.07	265	6,427.88	-6162.88	No	44.10	13:05	41/04	1,529.41	37.17	13:05	34.34
deletepool2	Step	2- Analyze	the capa	city with	Read I/C) Utilized	Avg 1,723.54		Max 6,677.17	Avg - 6,287.52			e Max Read	Avg 41.15	Max 1,936.19	Max 46.06		Avg 34.47
deletepool 1	Sep 14, 2011	OK	16.00	3.00	13.00	No	1,150.82	589	6,420.02	-5831.02	No I/O Ra &avail:	te with cap: ability	ability	39.24	1,585.84	42.87	12:00	30.37
	Sep 15, 2011	OK	16.00	3.00	13.00	No	0.00	589	0.04	588.96	Yes	81.80	00:20	63.39	0.00	0.00	00:20	0.00
deletepool1							Avg 575.41		Max 6,420.02	Avg - 2,621.03 /		Max 81.80		Avg 51.32	Max 1,585.84	Max 42.87		Avg 15.19
Cognos	Sep 14, 2011	Warning	92.00	79.90	11.25	No	1, 17	331	10.84	320.15	Yes	54.67	23:55	0.49	214.67	30.30	23:55	7.19
	Sep 15, 2011	Warning	92.00	79.90	11.25	No	1.53	331	51.63	279.37	No	21.17	24:00	0.76	103.33	31.37	24:00	5.26
Cognos							Avg 1.35		Max 51.63	Avg 299.76		Max 54.67	Treesence of the	Avg 0.62	Max 214.67	Max 31.37		Avg 6.22

Figure 15-99 Analyze the cognos MDisk Group for up-tiering

For the steps used, see Table 15-10), where the Mdisk Group **Cognos** is analyzed further to determine whether it is a suitable candidate for up-tiering.

Table 15-10 Steps to determine if the cognos MDisk Group is a canddate for up-tiering

Step	Parameters used to determine tiering	Observation	Up-tiering suggested	
1	MDisk Group Read I/O Rate Capability, MDisk Group Max Read I/O Rate, and MDisk Group Read I/O Rate Availability	The maximum read I/O rate varies from less than 10% in one instance . So, the availability rate of the MDisk group fluctuates between 80% to 90%	No	
2	MDisk Group Capacity Status, MDisk Group Capacity, MDisk Group Real Capacity, and MDisk Group Free Space	The capacity status is set to Warning. Consumed capacity is high (79.90 GB) and capability utilization is low fluctuating between 10% and 20%.	No	
3	MDisk Group Max Read Response Time, MDisk Group Max Read Response Time Threshold, MDisk Group Average Read Response Time, MDisk Group Max Write Response Time, and MDisk Group Max Write Response Time Threshold	Maximum and average read response times are high. Maximum and average write response times are high. The number of hours that the write response time exceeds the threshold is too high.	No	
Final re	commendation	MDisk Group Cognos is not over-utilized and its workloads are not candidates for up-tiering.		

Case Study 2: Identifying most active VDisks to reduce MDisk Group workload

VDisk based reports are helpful to report on the front-end status of your environment.

Overview

In this example we generate a report after using the following drill through navigation path:

- 1. MDisk Groups Workload Activity Report
- 2. MDisk Group VDisk Workload Activity Details Report
- 3. VDisk Details Report

Steps for analyzing VDisks

Follow these steps:

1. Use the most active or least active MDisk Group chart to select a MDisk Group to drill through to a VDisk (Figure 15-100).



Figure 15-100 Workload Activity Report - MDisk Groups

2. After clicking the **Cognos** MDisk Group (red bar), the VDisk Workload Activity Details Report opens and displays the 5 most active and the 5 least active VDisks in MDisk Group **Cognos** (Figure 15-101).



Figure 15-101 VDisk Activity Details report for Cognos MDisk Group

3. The most active VDisk in the **Cognos** MDisk Group is **tvt**. See Figure 15-102 for an analysis of the VDisk **tvt**.

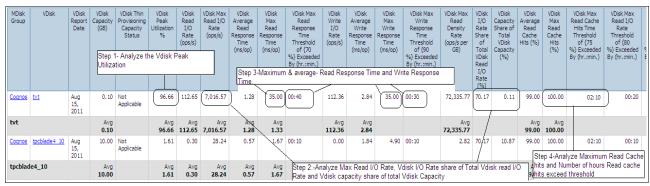


Figure 15-102 Analyze the tvt VDisk

Complete the steps in Table 15-11 to determine whether the VDisk requires up-tiering or down-tiering.

Table 15-11 Steps to determine whether VDisk tvt requires up-tiering or down-tiering

Step	Parameters used to determine tiering	Observation	Up-tiering suggested		
1	VDisk Peak Utilization	VDisk peak utilization is high.	Yes		
2	VDisk Max Read I/O Rate, VDisk I/O Rate Share of Total VDisk Read I/O	VDisk Max Read I/O Rate is very high. VDisk I/O Rate share of Total Vdisk read I/O Rate is high and VDisk capacity share of total VDisk Capacity is low.	Yes		
3	VDisk Max Read Response Time, VDisk Max Read Response Time Threshold, VDisk Max Write Response Time, and VDisk Max Write Response Time Threshold	Read and write response times are high and the number of hours that the read and write response time thresholds are exceeded are high.	Yes		
4	VDisk Max Read Cache Hits and VDisk Max Read Cache Hits Time Threshold	The maximum read cache hits and the number of hours that the maximum read cache hit threshold is exceeded is high.	Yes		
Final re	commendation	Migrating the VDisk tvt to another MDisk Group with adequate capacity would reduce the MDisk Group utilization. It will cool off the Cognos Mdisk Group.			

15.8.6 Considerations

This section describes considerations pertaining to the following topics.

Performance Monitor

A performance monitor must be scheduled to run for the entire time period for which you want to capture Storage Tier Report information about the IBM System Storage SAN Volume Controller or IBM Storwize V7000 device.

For details on setting up and running performance monitors, see the section "Planning for performance management" in Chapter 1 of the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide*, SC27-2337.

Tip: Instead of using a "continue indefinitely" Performance Monitor, we suggest that you run the Performance Monitor 24 hours a day and restart it again. This will ensure that all changes in the device configuration and also the job alerts are noted and included in the performance monitoring job.

The reason that we now can host a 24 hour performance monitor, instead of 23 hours, is that the Tivoli Storage Productivity Center for Disk code was enhanced in 4.2.1 to support the shutdown and restart without having the long delays seen in prior versions.

Reports based on Storage Resource Groups

If you plan to create reports based on Storage Resource Groups (SRGs), you must create the SRGs up front, and they must contain volumes, not subsystems, as seen in Figure 15-103.

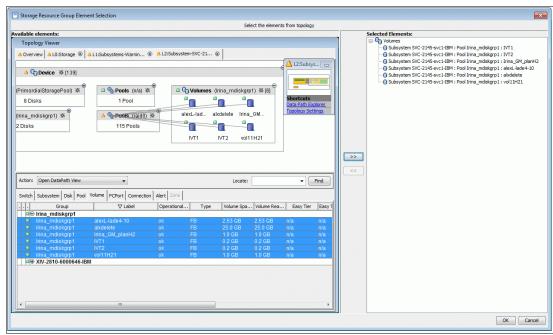


Figure 15-103 Creating SRG containing Volumes, not Subsystems

If your environment aligns to it, we suggest that you create around ten SRGs. This would create the most meaningful output from the bar charts. Creating less then ten SRGs will not affect reporting functionality in any way but could limit the meaningfulness of the bar chart reports as seen in Figure 15-104. Creating more than ten SRGs is perfectly fine.

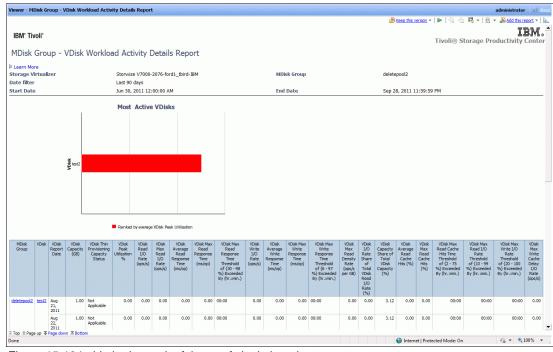


Figure 15-104 Limited meaningfulness of single bar chart reports

For your reports to pick up your changed SRG environment and to show data at all, you need to create performance monitors for the selected subsystems that you want to monitor.

To create performance monitors, follow the guidance in the section "Planning for performance management" in Chapter 1 of the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide*, SC27-2337.

Virtualizer boundaries

The current release of TCR only allows report creation up to a virtualizer boundary. Even if you add multiple volumes from different virtualizers into a single SRG, the report will only show output based on one single virtualizer. So the granularity of reports will be VDisks, SRGs, MDisk Groups, or Virtualizer.

Navigating through the reports

We strongly suggest that you *do not* use the web browser's **back** button. At any time, you are able to cancel out of a report creation process by using the **Cancel** button on the bottom of the page. If you are already within a report, you can always use the **back** button on the top right corner of your report as seen in Figure 15-105. If there is no **back** button on your current page, you are within a part of the report that has been opened in a new browser window. Feel free to close it at any time.

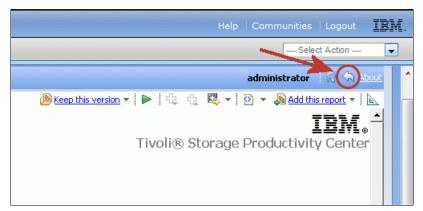


Figure 15-105 Back button within reports

Creating your own reports

With the current release of TCR, creating your own reports is not supported.

Configuring e-Mail part of TCR Reports

Configuration of the e-Mail part of the TCR Reports feature is done with the IBM Cognos Configuration Utility. To start the IBM Cognos Configuration utility, select $Start \rightarrow AII$ Programs \rightarrow Tivoli Common Reporting 2.1.1 \rightarrow IBM Cognos Configuration.

The IBM Cognos Configuration Utility starts and welcomes you with the panel shown in Figure 15-106.

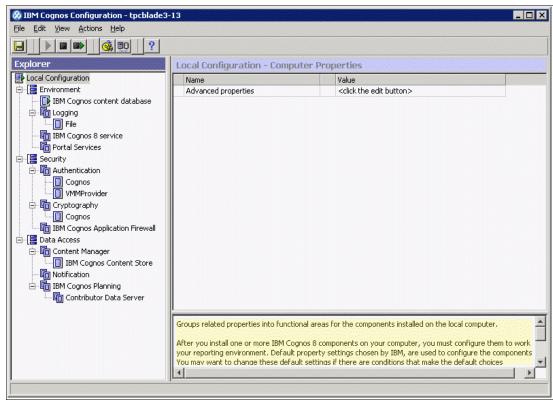


Figure 15-106 IBM Cognos Configuration Utility Welcome panel

Click **Notification** within the *Data Access* section of the left hand side Explorer. The Component Properties of the e-Mail feature will open as shown in Figure 15-107. Make the changes needed to enable emailing for your environment by entering SMTP mail server, Account and password, and Default sender information.

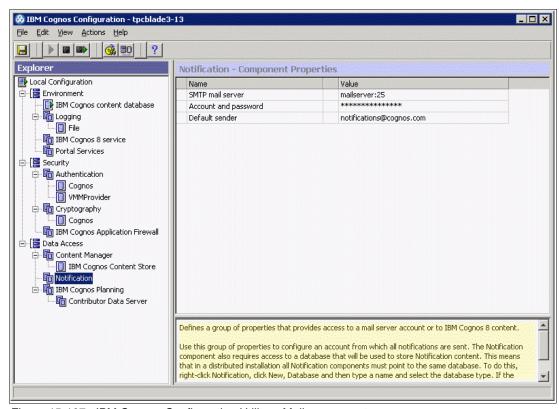


Figure 15-107 IBM Cognos Configuration Utility e-Mail component

After updating the configuration, you need to save your configuration and restart TCR to enable your changes. Click the **Save configuration** button or choose **Save** from the *File* menu to save your configuration. Your configuration will be saved as shown in Figure 15-108.

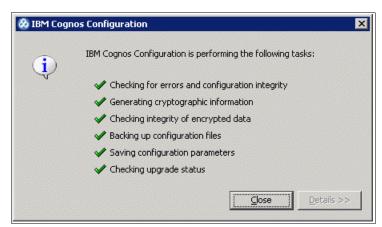


Figure 15-108 IBM Cognos Configuration Utility configuration saving

Afterwards, use the **Restart** button or choose **Restart** from the *Actions* menu to restart the service. TCR restarts as shown in Figure 15-109.

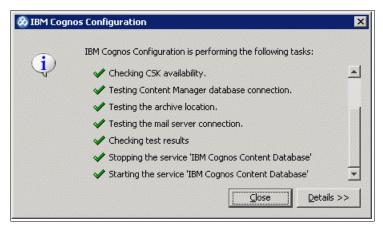


Figure 15-109 IBM Cognos Configuration Utility service restart

With these changes, you can now receive reports sent to you via e-Mail from TCR. Exit the IBM Cognos Configuration utility after you successfully completed your configuration changes.

Scheduling of TCR reports

TCR Reporting comes with a complete scheduling infrastructure to automatically create reports following schedules you can define. To create a report schedule, click the **Schedule** icon on the right hand side of every report that is available as shown in Figure 15-110.

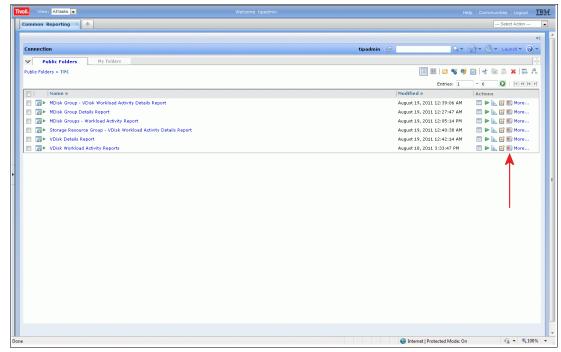


Figure 15-110 TCR Reports Scheduling Icon

In the Schedule panel shown in Figure 15-111 you can enter the details of your schedule. Choose the **Priority** of your schedule and enter Start and End date and time. You have the option to omit the End date for continuously running the schedule. Specify a frequency for your schedule on a **by day**, **by week**, **by month**, or **by year** granularity. Click **OK** to create your schedule.

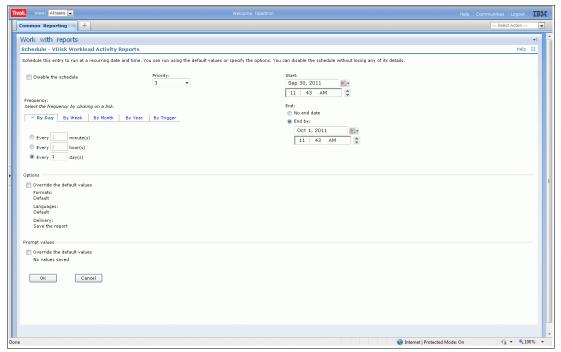


Figure 15-111 TCR Scheduling configuration

After the successful schedule creation, you can edit your schedule or delete it by clicking the **More...** button from the reports overview shown in Figure 15-112.

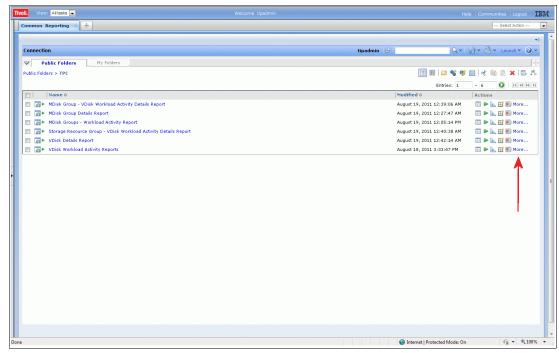


Figure 15-112 TCR Reports More... icon

You will be presented with the options shown in Figure 15-113. Choose **Modify the schedule** or **Remove the schedule** to work with your created schedule.

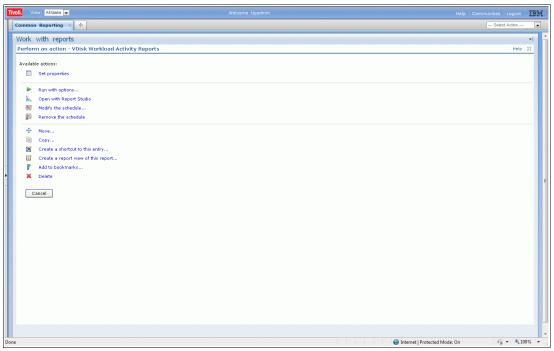


Figure 15-113 TCR Reports More... options

15.8.7 Migrating from the STAR tool

For those familiar with the Storage Tiering Activity Reporter (STAR) service offering from IBM and already collected data for tiering analysis using this tool, plan accordingly for transitioning to the Tivoli Storage Productivity Center Tiering Reports as follows:

- ▶ MDisk Groups using the "STAR" naming convention can be migrated to automatically set the back-end subsystem pool information.
- Tivoli Storage Productivity Center Tiering Reports will not work with migrated performance data because new metrics are collected in Tivoli Storage Productivity Center 4.2.2 that are required for the reports to run.
- STAR report configurations cannot be migrated to Tivoli Storage Productivity Center Storage Tier reports.
- Tivoli Storage Productivity Center Storage Tier reports do not provide all the same data that STAR reports provide.

15.8.8 Supported environments

The Tivoli Storage Productivity Center V4.2.2 Storage Tier reports support only *homogeneous* storage pools (HDD *or* SSD) in the following environments:

- Internal to SVC or Storwize V7000
- External to SVC or Storwize V7000

15.8.9 Non-supported environments

The Tivoli Storage Productivity Center V4.2.2 Tier Reports do not support or apply to the following environments:

- ► Hybrid pools:
 - A pool made up of storage from multiple back-end storage pools or subsystems
 - A pool with SSD and HDD at the same time
- ► EasyTier pools:
 - Tiering is done automatically by SVC and Storwize V7000.
 - Therefore, Tiering Reports do not apply.



Tivoli Storage Productivity Center database backup on Windows

This chapter explains how to plan for backing up and restoring the Tivoli Storage Productivity Center database that resides in DB2 on the Tivoli Storage Productivity Center server in a Windows environment.

The chapter covers both offline backup (cold backup) and online backup (hot backup) of the database along with the merits of each type.

The Tivoli Storage Productivity Center product does not provide any extra backup and recovery tools in addition to those tools already provided with the DB2 product.

This chapter is not intended to be a comprehensive guide to all functions of backup and recovery built into DB2. For details about this subject, see the *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441.

16.1 Before you start

We suggest that you read Chapter 17, "Tivoli Storage Productivity Center database considerations" on page 687 before you continue with this chapter. To improve the performance of your DB2 instance, tune your database settings by following the instructions that are provided in this section, and review 17.4.3, "Database backup method considerations" on page 711.

You can also estimate the storage requirements for the Tivoli Storage Productivity Center repository database, and plan the sizing of your backup volumes accordingly.

16.2 Common backup setup steps

Complete the following setup steps for both file system and Tivoli Storage Manager backups:

 Configure the DB2 history file to keep the number of backup versions that you want to retain. Your organization might already have a policy for how many versions you need to keep.

Change the DB2 num_db_backups parameter and set the value to the number of backup versions that you require. You must also set the rec_his_retentn parameter to a value of -1. By setting this value to -1, the rec_his_retentn parameter follows the value set in the num_db_backups parameter.

Important: For this change to take effect, Tivoli Storage Productivity Center service must be stopped and started. This restarting does not necessarily need to happen directly after you change the parameter.

2. Start a DB2 command-line processor window (Figure 16-1).

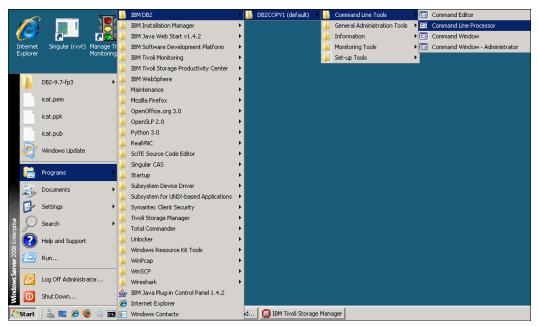


Figure 16-1 Launch the DB2 command line processor

3. A command-line processor window opens (Figure 16-2).

Figure 16-2 DB2 command line processor

4. Example 16-1 shows how to set the num_db_backups value to 4 versions, and set rec_his_retentn to -1 for both the Tivoli Storage Productivity Center database backups.
Issue the commands at the db2 => prompt in the command-line processor window.

Example 16-1 DB2 commands to configure how many backup versions to keep

```
connect to TPCDB
update db cfg using num_db_backups 4
update db cfg using rec_his_retentn -1
disconnect TPCDB
exit
```

Important: When you set new values for num_db_backups and rec_his_retentn, the new values are not effective until you stop all database connections.

Restart Tivoli Storage Productivity Center to make the changes effective. You can either reboot the server, or alternatively stop and start the services.

If you stop and start the services, use either the Windows Services interface or open a command prompt window and issue the commands in Example 16-2. (Use the **net start** command to obtain a list of active services in case your are using separate versions.) This process applies to Windows servers only.

Example 16-2 Windows commands to stop and start Tivoli Storage Productivity Center services

```
net stop "IBM Tivoli Storage Productivity Center - Data Server"
net stop "IBM WebSphere Application Server V6.1 - DeviceServer"

net start "IBM Tivoli Storage Productivity Center - Data Server"
net start "IBM WebSphere Application Server V6.1 - DeviceServer"
```

16.3 Offline backup to file system setup steps

This section describes how to set up offline backup for the Tivoli Storage Productivity Center server database to flat files in a file system. Because the offline backup method is the default method for Tivoli Storage Productivity Center, there is little DB2 configuration needed before you can perform a backup.

Important: Ensure that you perform the steps in 16.2, "Common backup setup steps" on page 650 and these steps.

The steps are as follows:

Choose a location to use for the DB2 backup output. Choose a directory that has enough
free space to hold the number of backups that you plan to retain. It is best to use a
separate file system rather than the file system that contains the DB2 database.

You can use a location that is a remotely mounted CIFS or NFS so that the backup data is secured to another server, perhaps at another location in your organization. This example uses D:\TPC_database_backups location.

Important: DB2 does not create this directory for you. Create this directory before you attempt a backup.

2. Create a batch script to control the backup process.

Two files are used:

- C:\scripts\TPC backup offline file.bat

This file (shown in Example 16-3) runs the backup.

Example 16-3 File C:\scripts\TPC_backup_offline_file.bat

```
@echo on
@REM This is a sample backup script
@REM To backup TPC offline
QREM To disk filesystems
@REM Stopping Tivoli Storage Productivity Center services
net stop "IBM Tivoli Storage Productivity Center - Data Server"
net stop "IBM WebSphere Application Server V6.1 - DeviceServer"
@REM Starting backup of the DB2 database
C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tv force application all
C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tvf
C:\scripts\database list offline file.txt
QREM Restarting Tivoli Storage Productivity Center services
net start "IBM Tivoli Storage Productivity Center - Data Server"
net start "IBM WebSphere Application Server V6.1 - DeviceServer"
@REM Offline backup process complete
@REM -----
```

- C:\scripts\database_list_offline_file.txt

This file (shown in Example 16-4) is a DB2 scripted list of databases to back up.

Example 16-4 File C:\scripts\database_list_offline_file.txt

backup database TPCDB to "D:\TPC_database_backups" without prompting;

See 16.7.1, "Performing an offline backup to a file system" on page 665 to run an offline backup.

16.4 Offline backup to Tivoli Storage Manager setup steps

This section describes the steps necessary to set up an offline backup of the Tivoli Storage Productivity Center server database to a Tivoli Storage Manager server. The backup to Tivoli Storage Manager is a little more complex to set up but does not require you to set aside large amounts of local disk space for backup versions on the Tivoli Storage Productivity Center server.

In this section, we assume the following requirements:

- You have a basic working knowledge of Tivoli Storage Manager.
- An operational Tivoli Storage Manager server already exists to which you can send backup data.
- ➤ Your Tivoli Storage Manager administrator has defined storage, which will receive the backups, to the policies.
- You have already installed a Tivoli Storage Manager Backup-Archive client on the Tivoli Storage Productivity Center server, and you have configured it to do standard file backups.
- You have installed the Tivoli Storage Manager API Client on the Tivoli Storage Productivity Center server.
- You used default installation paths for Tivoli Storage Manager.

Stop and reboot: You must stop Tivoli Storage Productivity Center and DB2 as part of this configuration process. Reboot the Tivoli Storage Productivity Center server to complete the configuration process, because this process also adds operating system environment variables. Plan this exercise at a time when you can reboot the Tivoli Storage Productivity Center server.

Use the following steps, described in this section, to configure DB2 to Tivoli Storage Manager integration:

- 1. "Adding new variables to Windows" on page 653
- 2. "Configuring Tivoli Storage Manager option file and password" on page 655
- 3. "Rebooting the Tivoli Storage Productivity Center server" on page 656
- 4. "Creating an offline backup to Tivoli Storage Manager script" on page 657

16.4.1 Adding new variables to Windows

Table 16-1 shows a list of Tivoli Storage Manager API environment variables to add to Windows. The listed values assume a default installation of Tivoli Storage Manager on the Tivoli Storage Productivity Center server.

Table 16-1 System environment variables

Environment variable name	Value
DSMI_DIR	C:\Program Files\Tivoli\TSM\baclient
DSMI_CONFIG	C:\Program Files\Tivoli\TSM\baclient\dsm.opt
DSMI_LOG	C:\tsm

The steps to add new variables to Windows are as follows:

1. In the Windows System Properties panel (Figure 16-3), click **Environment Variables** to proceed to the next step.

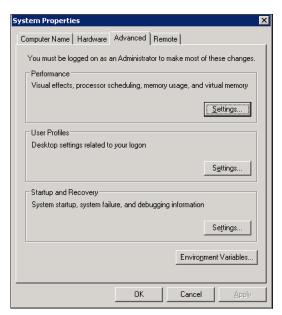


Figure 16-3 Windows System Properties

2. In the Environment Variables panel (Figure 16-4), click New.

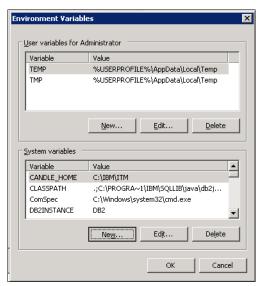


Figure 16-4 Windows Environment Variables

3. Add all three new system variables (as listed in Table 16-1 on page 654). Repeat the process for each variable. See Figure 16-5.



Figure 16-5 Adding a New System Variable

16.4.2 Configuring Tivoli Storage Manager option file and password

This section describes the steps necessary to configure the Tivoli Storage Manager option file dsm.opt, and then set the Tivoli Storage Manager password so that the DB2 backup process can communicate with the Tivoli Storage Manager API.

Important: At this stage, your Tivoli Storage Manager client must be already registered with a Tivoli Storage Manager server. If the Tivoli Storage Manager server is accepting open registrations, simply by starting the Tivoli Storage Manager client GUI or command line, you are asked for a password to register your client. If the Tivoli Storage Manager server is using closed registration, you will need the Tivoli Storage Manager administrator to register your client.

The steps are as follows:

- Edit the dsm.opt file, which is located in the following directory by default:
 C:\Program Files\Tivoli\TSM\baclient
- 2. Make sure that the client option PASSWORDACCESS is set to GENERATE, as shown in Figure 16-6, and save the file.

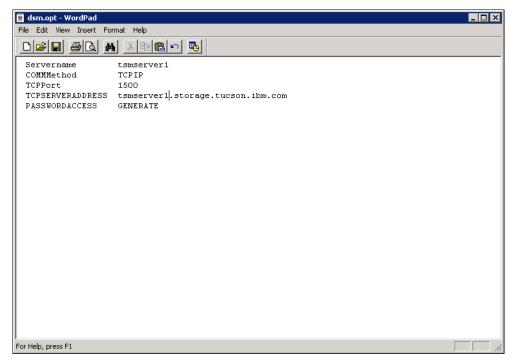


Figure 16-6 Example of a dsm.opt file

- 3. Set the Tivoli Storage Manager password so that DB2 can authenticate with the Tivoli Storage Manager server when DB2 performs a backup or restore operation:
 - a. Open a Windows Command prompt window.
 - b. Change to the location to run the dsmapipw.exe file. Run the dsmapipw command as shown in Figure 16-7.
 - c. Enter the current and new Tivoli Storage Manager password. You can reuse the existing Tivoli Storage Manager password.

Important: You must run the **dsmapipw** command even if you do not intend to change the Tivoli Storage Manager password. Running this command registers the password with the Tivoli Storage Manager API. Registering this password in the setup phase means that a DB2 operator can perform backup and restore operations without needing to know the Tivoli Storage Manager client password. If a Tivoli Storage Manager administrator changes or resets the Tivoli Storage Manager password, you need to run the **dsmapipw** command again.

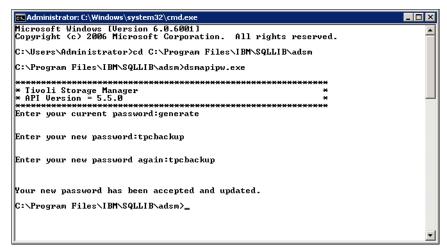


Figure 16-7 Running the dsmapipw command

Secure area: The dsmapipw command displays both the old and new passwords on the window in plain text. Ensure that you perform this task in a secure area to prevent password exposure.

16.4.3 Rebooting the Tivoli Storage Productivity Center server

Now that you have completed the configuration steps, reboot the Tivoli Storage Productivity Center server to ensure that the environment variables are picked up by DB2.

16.4.4 Creating an offline backup to Tivoli Storage Manager script

We based the script in Example 16-5 on a Tivoli Storage Productivity Center installation on Windows.

Create two files:

► The first file is the script that you run (Example 16-5):

```
C:\scripts\TPC backup offline tsm.bat
```

Example 16-5 File C:\scripts\TPC_backup_offline_tsm.bat

```
echo on
REM This is a sample backup script
REM To backup TPC offline
REM To Tivoli Storage Manager
REM Stopping Tivoli Storage Productivity Center services
net stop "IBM Tivoli Storage Productivity Center - Data Server"
net stop "IBM WebSphere Application Server V6.1 - DeviceServer"
REM Starting backup of the DB2 database
REM -----
C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tv force application all
C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tvf
C:\scripts\database list offline tsm.txt
REM Restarting Tivoli Storage Productivity Center services
net start "IBM Tivoli Storage Productivity Center - Data Server"
net start "IBM WebSphere Application Server V6.1 - DeviceServer"
REM Offline backup process complete
RFM -----
```

► The second file is the DB2 scripted list of databases to back up (Example 16-6):

```
C:\scripts\database_list_offline_tsm.txt
```

Example 16-6 File C:\scripts\database_list_offline_tsm.txt

backup database TPCDB use tsm without prompting;

16.5 Online backup to Tivoli Storage Manager setup steps

This section describes the steps for configuring the Tivoli Storage Productivity Center database to enable for online backup to Tivoli Storage Manager. The significant difference between online and offline backup is the need to enable archive logging on the database. As we describe in "Database backup method considerations" on page 711, operating in this mode provides many backup and recovery benefits at the expense of increased complexity in the database operation.

Important: You must stop Tivoli Storage Productivity Center to perform these tasks.

- ▶ DB2 requires a full backup of each database before you can start the Tivoli Storage Productivity Center database again after these reconfiguration steps. We include the instructions to perform a full backup of the database. Allow time in your outage planning for the backup to complete.
- ► Also, complete the steps in 16.2, "Common backup setup steps" on page 650 to set the number of backup versions that you want to retain in the history file.

Be sure you consider the advantages and disadvantages of archive logging before you continue with this setup. For full details of DB2 logging methods, see the DB2 product manuals. Also see *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441, for detailed information about this subject.

Set up and test DB2 to Tivoli Storage Manager integration before you attempt this section. Use 16.4, "Offline backup to Tivoli Storage Manager setup steps" on page 653. When you are satisfied that DB2 is communicating with Tivoli Storage Manager and you have performed at least one successful offline backup, return to this section.

16.5.1 DB2 parameter changes for archive logging to Tivoli Storage Manager

To set up archive logging to Tivoli Storage Manager, complete the following tasks:

1. You must make a number of parameter choices for the configuration of archive logging as shown in Table 16-2. These parameters determine where DB2 keeps its log files, the number of log files, and the size of the log files.

Tabla	100	DB2 parame	+
ianie	10-2	TIBZ narame	iers

DB2 parameter	Example value	Comment
Primary log path	C:\DB2_active_logs	This is the location where DB2 keeps the current logs for the database. For best performance, place these logs on a separate volume than the volume that holds the data.
Failed log path	D:\DB2_failed_log	This is the location where DB2 put log files if the archive process fails. This can happen if Tivoli Storage Manager is down or unreachable when DB2 tries to send a log file to Tivoli Storage Manager.

2. Stop Tivoli Storage Productivity Center by using the commands in Example 16-7. You can also perform this task through the Windows Services interface.

Example 16-7 Windows commands to stop Tivoli Storage Productivity Center

net stop "IBM Tivoli Storage Productivity Center - Data Server"

3. Launch a DB2 command-line processor, as in Figure 16-8, and issue the commands shown in Example 16-8.

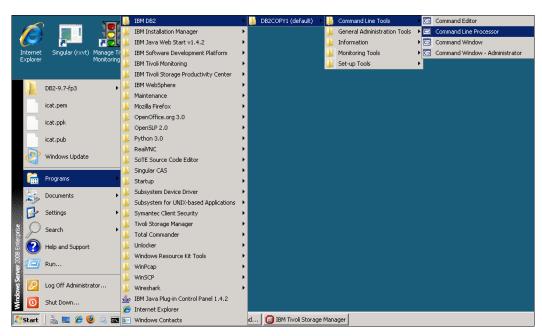


Figure 16-8 Launch a DB2 command line processor

A command-line processor opens (Figure 16-9).

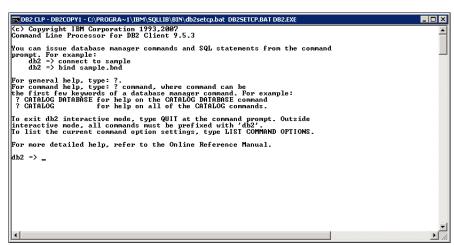


Figure 16-9 DB2 command line processor

4. Issue the commands shown in Example 16-8 in the command-line processor window. Substitute your chosen values for the parameters that form part of the UPDATE DB CFG command (see Table 16-2 on page 658). The final command performs an offline backup of the database.

Important: The database backup is required after this reconfiguration, and the DB2 database will not open again until the database backup is completed.

```
CONNECT TO TPCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR TPCDB USING logarchmeth1 TSM failarchpath "D:\DB2_failed_logs" newlogpath C:\DB2_active_logs\TPCD

BACKUP DATABASE TPCDB USE TSM

QUIT
```

When the database backups is complete, restart Tivoli Storage Productivity Center. Either use the Windows Services interface or issue the commands shown in Example 16-9 in a command window.

```
Example 16-9 Start Tivoli Storage Productivity Center
```

```
net start "IBM Tivoli Storage Productivity Center - Data Server"
net start "IBM WebSphere Application Server V6.1 - DeviceServer"
```

16.5.2 Creating an online backup script for Tivoli Storage Manager

We based this example on a Tivoli Storage Productivity Center installation on Windows.

Create two files:

► The script (Example 16-10) that you run to start the backup:

C:\scripts\TPC_backup_online_tsm.bat

Example 16-10 File C:\scripts\TPC_backup_online_tsm.bat

```
echo on

REM This is a sample backup script
REM To backup TPC online
REM To Tivoli Storage Manager

REM Starting backup of the DB2 database
REM ------

C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tvf
C:\scripts\database_list_online_tsm.txt

REM Offline backup process complete
REM ------
```

► The DB2 scripted list (Example 16-11) of databases to back up:

C:\scripts\database_list_online_file.txt

Example 16-11 File C:\scripts\database_list_online_tsm.txt

backup database TPCDB online use tsm without prompting;

16.6 Online backup to a file system setup steps

Performing online backups to a file system requires you to set up archive logging to a file system also. When operating with this method, DB2 does not clean up archive log files that are old and no longer necessary. Therefore, you must put processes in place to clean up old log files after a specific amount of time to prevent the system from filling up. You must also plan for this amount of space. Over a number of weeks, the log space that is required for a Tivoli Storage Productivity Center database can become many times larger than the database.

To be able to restore an online DB2 database taken two weeks ago, for example, you must have log files going back to that same date that you can restore. An online DB2 database backup is not stand-alone, because you cannot restore the online DB2 database backup without at least some logs for it to roll forward to a consistent state.

Important:

- ► Although switching between a backup destination from online to file system and online to Tivoli Storage Manager is straightforward, switching the logging path is not as easy. To switch the logging from Tivoli Storage Manager to a file system requires a stop and a start of the database and, therefore, a restart of the Tivoli Storage Productivity Center services.
- ► Choose either a Tivoli Storage Manager backup or a file system backup, and stay with that specific method.

16.6.1 Setting up DB2 archive logging to a file system

Set up DB2 archive logging to a file system by using the following steps:

1. To configure archive logging, you must make several parameter choices. The parameters in Table 16-3 determine where DB2 keeps its log files, how many log files to keep, and the size of the log files.

Table 16-3 DB2 parameters for archive logging to a file system

DB2 parameter	Example value	Comment
Primary log path	C:\DB2_active_logs	The location where DB2 will keep the current logs for the database. For best performance, place the logs on a separate volume from the data.
Archive log path	C:\DB2_archive_logs\TPCDB	The location where DB2 will archive log files for the TPCDB database.
Failed log path	D:\DB2_failed_log	The location where DB2 will put log files if the archive process fails, which can happen if the file system for the primary logs fills up. Choose a location that is <i>not</i> on the same file system as the archive logs.

2. Choose a file system path to store the DB2 database backups. See Table 16-4.

Table 16-4 File system location for database backups

Database backup path
D:\TPC_database_backups

3. Stop Tivoli Storage Productivity Center by using the commands in Example 16-12. You can also perform this task through the Windows Services interface.

Example 16-12 Windows commands to stop Tivoli Storage Productivity Center

net stop "IBM Tivoli Storage Productivity Center - Data Server" net stop "IBM WebSphere Application Server V6.1 - DeviceServer"

4. Launch a DB2 command-line processor (see Figure 16-10) to issue the commands.

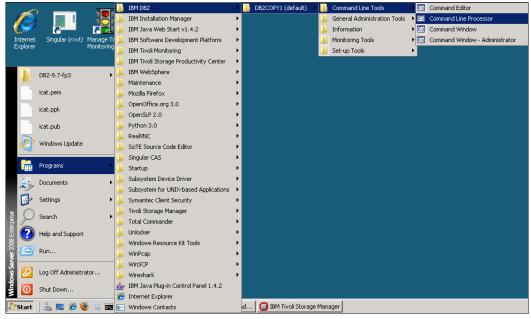


Figure 16-10 Launch a DB2 command line processor

A DB2 command-line processor opens (Figure 16-11).

Figure 16-11 DB2 command line processor

Issue the commands shown in Example 16-13 on page 663 in the command window.
 Substitute your values for the parameters that form part of the UPDATE DB CFG command. See Table 16-3 on page 662. Note that the final command performs an offline backup of the database.

Important: The offline backup of the database is required after the reconfiguration; the DB2 database does not open until the backup is complete.

Example 16-13 DB2 command to configure archive logging to a file system

CONNECT TO TPCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR TPCDB USING logarchmeth1 "DISK:C:\DB2_archive_logs"
failarchpath "D:\DB2_failed_logs" newlogpath C:\DB2_active_logs\TPCD

BACKUP DATABASE TPCDB TO "D:\TPC_database_backups"

When both database backups complete, you can restart Tivoli Storage Productivity
Center. Either use the Windows Services interface or issue the commands shown in
Example 16-14 in a command window.

Example 16-14 Start Tivoli Storage Productivity Center

net start "IBM Tivoli Storage Productivity Center - Data Server" net start "IBM WebSphere Application Server V6.1 - DeviceServer"

16.6.2 Creating an online backup script to a file system

Create files to control the backup process for online backup to file system output. We based this example on a Tivoli Storage Productivity Center installation on Windows.

Create two files:

► The script file (Example 16-15) that you run to start the backup:

```
C:\scripts\TPC_backup_online_file.bat
```

Example 16-15 File C:\scripts\TPC_backup_online_file.bat

```
REM This is a sample backup script
REM To backup TPC online
REM To filesystem

REM Starting backup of the DB2 database
REM -----

C:\PROGRA~1\IBM\SQLLIB\BIN\db2cmd.exe /c /w /i db2 -tvf
C:\scripts\database_list_online_file.txt

REM Offline backup process complete
REM ------
```

► The DB2 scripted list of databases (Example 16-16) to back up:

```
C:\scripts\database_list_online_file.txt
```

Example 16-16 File C:\scripts\database_list_online_file.txt

backup database TPCDB online to "D:\TPC_database_backups" without prompting;

16.7 Performing offline database backups

This section describes how to perform the offline backup of the Tivoli Storage Productivity Center database.

Running an offline DB2 database backup takes Tivoli Storage Productivity Center out of service for the period of the backup. This affects data collections from CIMOMs, and you might miss other infrastructure events.

Make sure that you understand the impact of stopping Tivoli Storage Productivity Center in your environment before proceeding. If your environment cannot tolerate regular stoppages of Tivoli Storage Productivity Center for a backup operation, consider configuring and using online backup.

16.7.1 Performing an offline backup to a file system

Important: You must complete the initial steps in 16.2, "Common backup setup steps" on page 650 before you can start offline backups.

To do an offline backup to a file system, run the <code>TPC_backup_offline_file.bat</code> script in a command window, as shown in Figure 16-12. This script stops the Tivoli Storage Productivity Center process, performs a backup of the DB2 database to the location specified in the scripts, and then restarts the services.

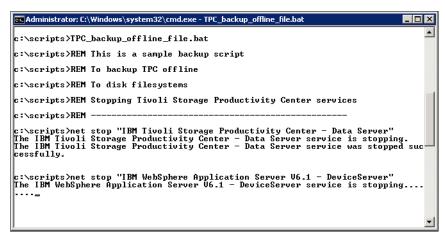


Figure 16-12 Running an offline backup to a file system

16.7.2 Performing an offline backup to Tivoli Storage Manager

Important: You must complete the initial setup steps described in 16.4, "Offline backup to Tivoli Storage Manager setup steps" on page 653 before you can start offline backups.

Running an offline DB2 database backup takes Tivoli Storage Productivity Center out of service for the period of the backup. Make sure it is acceptable to take Tivoli Storage Productivity Center out of service before you proceed.

To perform an offline backup to Tivoli Storage Manager, run the TPC_backup_offline_tsm.bat script in a command window as shown in Figure 16-13.

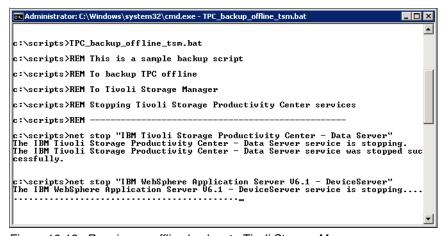


Figure 16-13 Running an offline backup to Tivoli Storage Manager

16.8 Performing online database backup

This section describes how to run online backups of the Tivoli Storage Productivity Center database. By running the backups online (or hot), it is not necessary to stop the Tivoli Storage Productivity Center service, which means that the backup process does not interrupt processes, such as performance data collection.

Important: Before you can perform online DB2 backup, you must configure Tivoli Storage Manager to DB2 communication and DB2 archive logging to Tivoli Storage Manager or DB2 to file system logging.

If you plan to use online backup to Tivoli Storage Manager, perform the setup steps in 16.4, "Offline backup to Tivoli Storage Manager setup steps" on page 653 and 16.5, "Online backup to Tivoli Storage Manager setup steps" on page 658.

16.8.1 Performing an online database backup to Tivoli Storage Manager

You can run an online backup of the Tivoli Storage Productivity Center database to Tivoli Storage Manager. You must already have run the appropriate setup steps.

To perform the backup, run the C:\scripts\TPC_backup_online_TSM.bat script in a command window as shown in Figure 16-14.

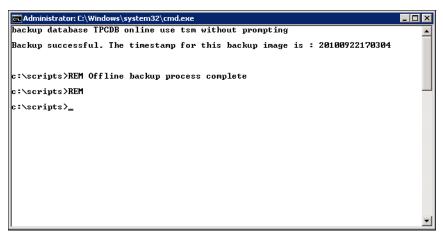


Figure 16-14 Running an online backup to Tivoli Storage Manager

16.8.2 Performing an online backup to a file system

You can run an online backup of the Tivoli Storage Productivity Center database to file system output files. You must have already completed the necessary setup steps detailed in 16.6, "Online backup to a file system setup steps" on page 661.

To do the backup, run the <code>C:\scripts\TPC_backup_online_file.bat</code> script in a command window as shown in Figure 16-15.

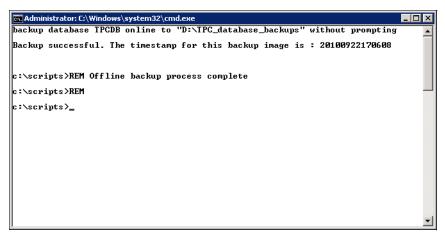


Figure 16-15 Performing an online backup to file system output

16.9 Other backup considerations

Apart from the DB2 database, there are a number of important files and directories to back up to preserve the state of a Tivoli Storage Productivity Center server installation.

Back up all files under the Tivoli Storage Productivity Center installation directory. In addition, the Tivoli GUID is stored in the Registry on Windows. Be sure that you also preserve this Tivoli GUID.

The important Tivoli Storage Productivity Center server directories to secure are as follows:

- TPC_Server_install_dir/config/
- ► TPC_Server_install_dir/data/config/
- TPC_Server_install_dir/device/conf/

These directories contain various configuration files for your installation. Saving these directories is important because they might be customized configurations and not the default configurations.

16.10 Managing database backup versions

This section is an overview of the maintenance processes for which you need to plan. You must maintain the number of Tivoli Storage Productivity Center database backup versions that you want on a file system or Tivoli Storage Manager because DB2 does not prune older versions automatically.

16.10.1 Managing backup versions for a file system

This section describes what you need to know to manage DB2 backups that were performed to disk.

Deleting unwanted backups

DB2 does not automatically manage the deletion of the unwanted database backups or archive logs from the file system. You need to create a maintenance plan to delete the old backups. If you plan to use online backup to a file system, you must create a plan to delete the old archive logs.

How DB2 organizes backups on a file system

When you perform a backup to a file system, you supply the backup script with a path to use. Version 9.1 of DB2 creates backup files with a file name containing a datestamp that details when the backup was taken.

Figure 16-16 shows the backup directory of the TPCDB database. This directory contains backups taken at separate times on 22 September 2010. DB2 timestamps all backups in this way; every time a backup is made, a new file is created:

- ► The first part of the file name starts as follows: TPCDB.0.DB2.NODE0000.CATN0000
- ► The last part of the file name consists of the date in *yyyyMMDD* format, as in this example: 20100922

Plan to delete old backup files to suit the requirements of your backup and recovery policy.

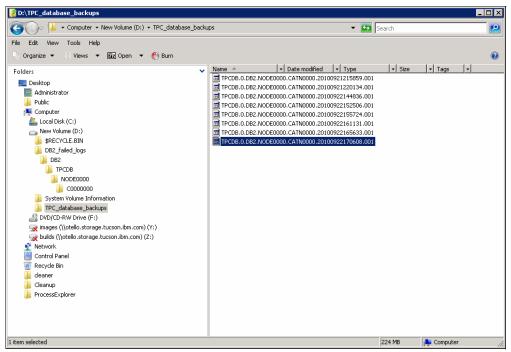


Figure 16-16 DB2 backup directory

16.10.2 Managing archive log files on a file system

If you plan to do online backups, you must configure DB2 to use archive logging. If you plan to do online backups to disk, you must also maintain the archive logs directory on a regular basis.

Figure 16-17 shows the directory structure for logging on to the TPCDB database. Over time, this directory fills up with logs. If your recovery policy is to keep backup versions for five days, you must keep logs in this directory for at least the same period of time, because you cannot restore an online backup without logs from the same date and time to make the recovery valid.

Notice that the directory that holds the logs is named C0000000, which is the log cycle number. If you restore the database, the cycle number increments by one and starts in C0000001 and so on. Ensure that any automated deletion process that you implement can handle this numbering.

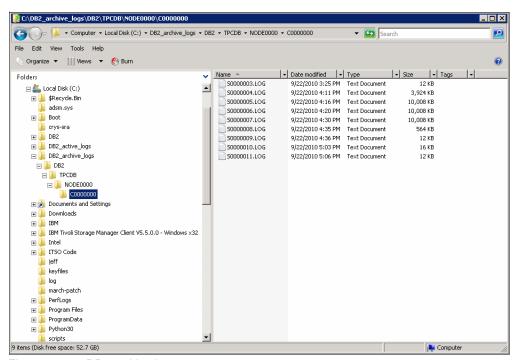


Figure 16-17 DB2 archive logs

16.10.3 Managing backup versions that you store in Tivoli Storage Manager

This section describes how to maintain, view, and delete backup data and archive logs that you have sent to Tivoli Storage Manager. DB2 does not automatically prune backup versions and log files from Tivoli Storage Manager. You need to use the db2adut1 tool to perform these housekeeping functions.

Tip: This section is not intended to be a comprehensive guide to the **db2adut1** tool. The intent here is to detail the commands that you likely will use on a regular basis to maintain the data that is held in Tivoli Storage Manager.

Purpose of the db2adutl command

The command line tool, db2adut1, communicates with Tivoli Storage Manager through its API interface. Use this tool to interrogate the backup and archive log data that is stored in Tivoli Storage Manager at any one time, verify that you no longer require old backups, and delete unnecessary old backups.

Reasons to use db2adutl

When DB2 stores a backup session in Tivoli Storage Manager, DB2 always stores the backup session with a unique file name, which is the timestamp when the backup was made. This means that these backup versions never get superseded by a new version with the same file name. The backup files remain "active" versions in Tivoli Storage Manager, and, therefore, Tivoli Storage Manager never deletes the backup versions. Use the <code>db2adut1</code> command to select unwanted backup versions and tell Tivoli Storage Manager to flag them as "inactive." In this way, Tivoli Storage Manager can then delete them over time, based on the standard policy rules that the Tivoli Storage Manager administrator sets.

You handle DB2 archive logs differently. They are stored in Tivoli Storage Manager as "archive" data, which means Tivoli Storage Manager retains them for a set period of time based on its policies. Although you can use **db2adut1** to explicitly remove DB2 archive logs, if Tivoli Storage Manager archive retention policy is set appropriately, this way is not necessary.

Important: Make sure that the Tivoli Storage Manager archive retention policy that you use to store the DB2 logs is set for a sufficient period of time to allow recovery of your oldest database backup. However, you also want to make sure that the policy for the retention period is not so long that it wastes storage space in Tivoli Storage Manager.

How to query backups held in Tivoli Storage Manager

Next, we explain how to query backups that are held in Tivoli Storage Manager.

Tip: You invoke **db2adut1** from a standard Windows command window. You have already set the path information for this command by the DB2 installation process. This command is normally in the SQLLIB\bin directory of DB2.

The db2adut1 commands to query database backup versions are as follows:

► The following command lists all the database versions and the logs that are held for all databases stored in Tivoli Storage Manager and TPCDB, in this case.

db2adut1 query

► The following command lists all database versions and logs for the TPCDB database. Note that the database name is case sensitive and is in capital letters.

db2adutl query database TPCDB

Figure 16-18 shows the sample output from this command. The output shows that two database backups are stored in Tivoli Storage Manager and six archive logs.

► The following command has a shorter output. It lists only the database backup versions and the archive logs.

db2adutl query full

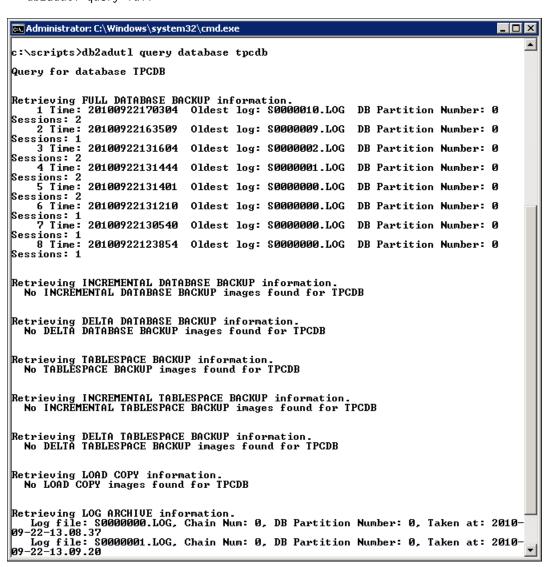


Figure 16-18 Sample output from a db2adutl query database TPCDB command

Deleting backup versions held in Tivoli Storage Manager

The commands and examples described next show how to delete database backup versions that are held in Tivoli Storage Manager:

► The following command deletes backup versions from Tivoli Storage Manager that are older than three days. This type of command is useful because you can easily script it to run each day to remove the next oldest backup.

```
db2adutl delete full older than 3 days
```

You can also specify a database name:

db2adutl delete full older than 3 days database TPCDB

Figure 16-19 is an example of running this command.

The following command deletes all backup versions from Tivoli Storage Manager, except for the last three versions. Again, this command is useful when scripting an automatic process.

```
db2adutl delete full keep 3
```

You can also specify a database name:

db2adut1 delete full keep 3 database TPCDB

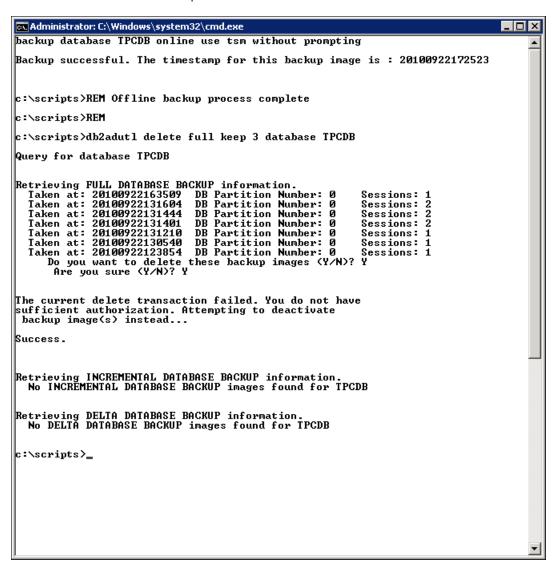


Figure 16-19 Example of a db2adutl delete full keep 3 database TPCDB command

Managing DB2 archive log files in Tivoli Storage Manager

The commands described in this section are examples of how to delete database archive logs from Tivoli Storage Manager.

You invoke the db2adut1 command from a standard Windows command window.

Important:

- ▶ Be careful when you delete archive log files. If you delete logs that are still needed for some of your backup versions, you render those backups *useless*.
- Archive logs only exist in Tivoli Storage Manager if you have configured archive logging so that online backup is possible.
- Ask the Tivoli Storage Manager administrator to configure Tivoli Storage Manager to delete the archive logs on a regular basis by configuring the Tivoli Storage Manager "archive copy group" that DB2 uses. Set a retention period that suits your needs. If you use a general purpose archive copy group, Tivoli Storage Manager might keep all archive logs for several years causing unnecessary usage of the storage in your Tivoli Storage Manager environment.

To delete archive logs, first query the Tivoli Storage Manager server to establish which logs you want to delete. Figure 16-18 on page 671 shows example output.

To query the Tivoli Storage Manager server for the TPCDB database, issue the following command:

db2adutl query database TPCDB

First, look at the "oldest" log number against the oldest backup version. After we delete several backups, as shown in Figure 16-19 on page 672, the oldest log is \$0000010.log.

Next, look at the list of log files from the same output to determine whether any earlier logs exist. If earlier logs do exist and you do not want to wait for Tivoli Storage Manager to expire them, use the following command to delete them. See Figure 16-20.

db2adutl delete logs between S0000001 and S0000004 database TPCDB

Tip: When you specify log numbers, add the **S** at the start of the number, but not the **.LOG** at the end.

```
Administrator: C:\Windows\system32\cmd.exe

c:\scripts\db2adut1 delete logs between $0000000 and $0000002 database TPCDB

Query for database TPCDB

Retrieving LOG ARCHIVE information.
Log file: $0000000.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2010-09-22-13.08.37
Do you want to delete this log image (Y/N)? Y
Are you sure (Y/N)? Y
Log file: $0000001.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2010-09-22-13.09.20
Do you want to delete this log image (Y/N)? Y
Are you sure (Y/N)? Y
Log file: $0000002.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2010-09-22-13.10.40
Do you want to delete this log image (Y/N)? Y
Are you sure (Y/N)? Y

C:\scripts\
c:\scripts\
c:\scripts\
c:\scripts\
c:\scripts\
```

Figure 16-20 Example command to delete DB2 archive logs

16.11 Verifying a backup file

To ensure that a backup file is valid and you are able to restore from it, you can use the db2adut1 command to check the integrity of a specific backup, entering the backup date stamp and database name as parameters:

db2adutl verify full taken at 20100922172357 db TPCDB

A verification process begins, as shown in Figure 16-21, and the following confirmation message is issued:

Image Verification Complete - successful.

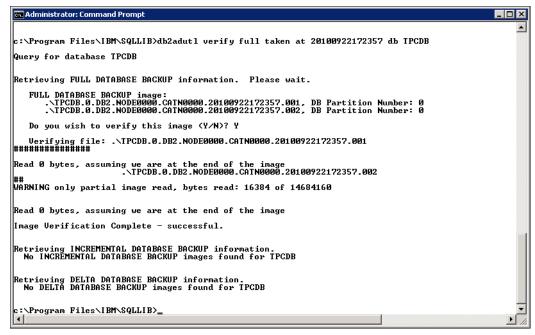


Figure 16-21 Performing a backup verification

If the verification fails, that backup is not usable and you will need to take a new one.

16.12 Restoring Tivoli Storage Productivity Center database

This section describes the steps to restore the DB2 repository database for Tivoli Storage Productivity Center. As with the backup process, restoring from an online backup is more complex than restoring from an offline backup.

Restoring from an offline backup is a simple point-in-time exercise. Because the database was stopped at the time of the offline backup, it is logically consistent and you can restore the data "as is." However, circular logging does not offer the ability to roll forward through database changes using the logs to recover to an exact point-in-time. Therefore, if you take a database backup on a 24-hour cycle, you lose updates to the Tivoli Storage Productivity Center repository that were made between these points.

When you configure archive logging, you can restore a backup and then roll forward through the logs to any point-in-time to minimize data loss. This way gives you an enhanced level of protection to the Tivoli Storage Productivity Center repository data at the expense of more complexity in the process. You cannot simply restore a backup taken online as is, because an online backup is not logically consistent in its own right. Following an online restore, some roll forward is necessary to bring the restored database to a consistent and usable state.

Finally, we do not intend for this section to be a comprehensive guide to the DB2 restore commands. We intend to give you the basic restore functions that you need to recover a database from both file system and Tivoli Storage Manager backups. For details about this subject, see *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441.

16.12.1 Restoring from offline backups

Restoring from an offline backup is the most simple type of restore. It brings the database back to the specific point-in-time that the backup was taken. You can then restart Tivoli Storage Productivity Center.

Overview of basic steps

Restoring an offline backup from a file system or Tivoli Storage Manager involves the following basic steps:

- 1. Stop the Tivoli Storage Productivity Center services if they are still running.
- 2. Choose the backup image from which to restore.
- 3. Restore the TPCDB database.
- 4. Restart the Tivoli Storage Productivity Center service.
- 5. Resolve potential agent issues after you restore. For more information, see 16.12.3, "Potential agent issues after the restore process" on page 685.

Stopping the Tivoli Storage Productivity Center services

Stop the Tivoli Storage Productivity Center services on Windows by using the commands in Example 16-17.

Example 16-17 Windows commands to stop Tivoli Storage Productivity Center

```
net stop "IBM Tivoli Storage Productivity Center - Data Server" net stop "IBM WebSphere Application Server V6.1 - DeviceServer"
```

Choosing the backup image to restore from file system

If the backup image that you require is stored in Tivoli Storage Manager, skip this section and follow the steps in "Choosing a backup image to restore from Tivoli Storage Manager" on page 676.

Use one of the following steps:

- Use Windows Explorer to look at the file system where you stored the backups and choose a backup image from which to restore
- ► Use the DB2 command list history backup all for TPCDB (in a DB2 command window) to see a list of the backup versions that are available.

Figure 16-22 shows that a backup image for the TPCDB database; the following file name is selected (which translates to 21 September 2010 at 21:58:59 timestamp):

TPCDB.O.DB2.NODE0000.CATN0000.20100921215859

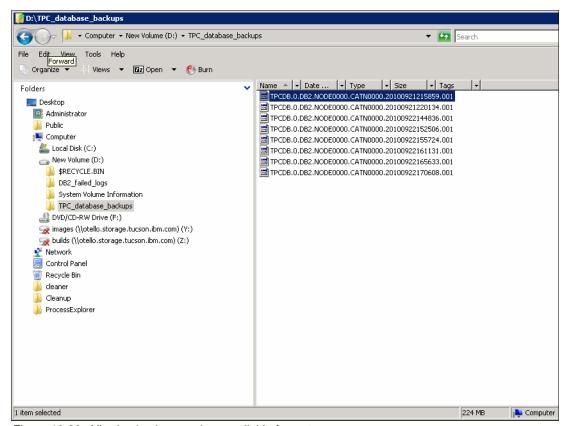


Figure 16-22 Viewing backup versions available for restore

In this example, the backup image timestamp is as follows:

20100921215859

You need this timestamp number for the next step, "Restoring the TPCDB database (offline)" on page 677.

Choosing a backup image to restore from Tivoli Storage Manager

If you have chosen a backup image from the file system, skip this step and move on to "Restoring the TPCDB database (offline)" on page 677.

To search for a backup image in Tivoli Storage Manager, use the db2adut1 command: db2adut1 query full database TPCDB

Figure 16-23 shows example output from the db2adut1 command for the TPCDB database.

```
Administrator: C:\Windows\system32\cmd.exe

c:\scripts>\db2adutl query full database TPCDB

Query for database TPCDB

Retrieving FULL DATABASE BACKUP information.
1 Time: 20100922172523 Oldest log: $0000013.LOG DB Partition Number: 0 Sessions: 2
2 Time: 20100922172357 Oldest log: $0000012.LOG DB Partition Number: 0 Sessions: 2
3 Time: 20100922170304 Oldest log: $0000010.LOG DB Partition Number: 0 Sessions: 2

Retrieving INCREMENTAL DATABASE BACKUP information.
No INCREMENTAL DATABASE BACKUP images found for TPCDB

Retrieving DELTA DATABASE BACKUP information.
No DELTA DATABASE BACKUP images found for TPCDB
```

Figure 16-23 Command db2adutl example to query backup versions available

In this example, the backup image timestamp is as follows:

20100922170304

You need this timestamp number for the next step, "Restoring the TPCDB database (offline)".

Restoring the TPCDB database (offline)

To restore the database, open a DB2 command-line processor window (Figure 16-24).

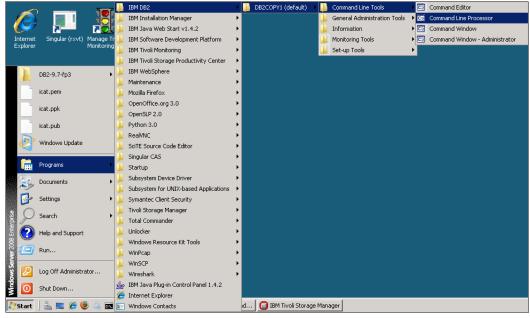


Figure 16-24 Launch a DB2 command line processor

A command window opens (Figure 16-25).

Figure 16-25 DB2 command line processor

To restore from file system backups, issue the commands (in Example 16-18) in the DB2 command window and use the timestamps that you selected.

Example 16-18 Restore command from file system backups

restore database TPCDB from "D:\TPC_database_backups" taken at 20100921215859

If you restore from Tivoli Storage Manager, use the commands shown in Example 16-19.

Example 16-19 Restore command from Tivoli Storage Manager backups

restore database TPCDB use TSM taken at 20100922170304

Figure 16-26 shows an example of the restore process dialog for the TPCDB database restore process from a file system.

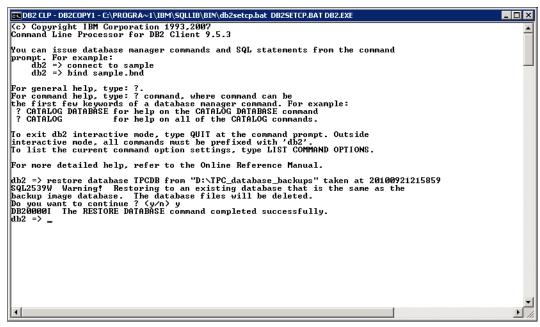


Figure 16-26 Example of offline restore of TPCDB from a file system

Restarting the Tivoli Storage Productivity Center services

After you restore the Tivoli Storage Productivity Center database, restart the Tivoli Storage Productivity Center services to bring the server back online. To do this on a Windows Tivoli Storage Productivity Center server, issue the commands, shown in Example 16-20, in a command window.

Example 16-20 Windows commands to start Tivoli Storage Productivity Center

```
net start "IBM Tivoli Storage Productivity Center - Data Server"
net start "IBM WebSphere Application Server V6.1 - DeviceServer"
```

16.12.2 Restoring from online backups

Restoring from an online backup can be more complex than restoring from an offline backup, because there are more choices about what to do after you restore the backup image.

You might restore to a backup image from a week ago, because you actually want your Tivoli Storage Productivity Center environment put back to that point. You might want to restore from the last known good backup and roll forward through the archive logs to get your Tivoli Storage Productivity Center database as close as possible to the point before the problem occurred that triggered the need to restore.

Overview of basic steps

Restoring an online backup from a file system or Tivoli Storage Manager involves the following basic steps:

- 1. Stop the Tivoli Storage Productivity Center services if they are not already stopped.
- Choose the backup image from which to restore.
- Restore the TPCDB database.

- 4. Roll forward the database.
- 5. Restart the Tivoli Storage Productivity Center services.
- 6. Resolve any new agent issues after you restore.

Stopping the Tivoli Storage Productivity Center services

Stop the Tivoli Storage Productivity Center services on Windows using the commands in Example 16-21. The services might already be stopped if something is not working correctly.

Example 16-21 Windows commands to stop Tivoli Storage Productivity Center

```
net stop "IBM Tivoli Storage Productivity Center - Data Server"
net stop "IBM WebSphere Application Server V6.1 - DeviceServer"
```

Choosing the backup image from which to restore

Choose a backup image from which to restore using the same process as you use for offline backups.

See "Choosing the backup image to restore from file system" on page 675 or "Choosing a backup image to restore from Tivoli Storage Manager" on page 676.

Restoring the TPCDB database (online)

The initial process of restoring a database that was taken online is the same as the offline process. However, when you complete the restoration, you are *not* ready to use the database. After you restore the backup, the database status is "Roll-Forward Pending" (Proceeding from this point is explained in "Rolling the database forward" on page 682).

To restore the database, open a DB2 command window as shown in Figure 16-27.

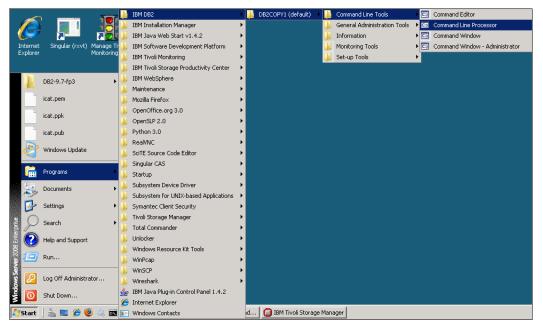


Figure 16-27 Launch a DB2 command line processor

A command window opens (Figure 16-28).

```
Co Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.5.3
You can issue database manager commands and SQL statements from the command prompt. For example:

db2 => connect to sample
db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ?.
For command help, type: ? CATALOG DATABASE command?
? CATALOG DATABASE for help on the CATALOG DATABASE command?
? CATALOG interactive mode, type QUIT at the command prompt. Outside interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => ____
```

Figure 16-28 DB2 command line processor

To restore the database from file system backups, issue the commands in Example 16-22 in the DB2 command-line processor using the timestamp that you have selected.

Example 16-22 Restore command from file system backups

restore database TPCDB from "D:\TPC_database_backups" taken at **20100924135535**

If you restore from Tivoli Storage Manager, use commands in Example 16-23.

Example 16-23 Restore command from Tivoli Storage Manager backups

restore database TPCDB use TSM taken at 20100924124117

Figure 16-29 shows an example of the restore process dialog for the TPCDB database restore from a file system.

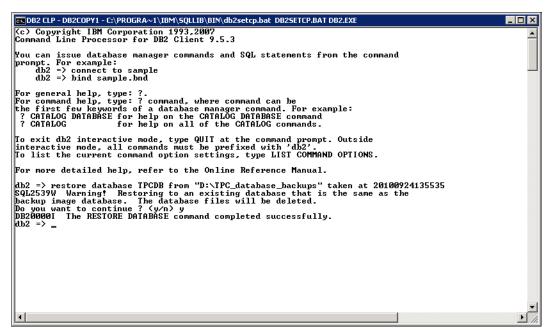


Figure 16-29 Example of online restore of TPCDB from a file system

Rolling the database forward

After the database restore processes complete, you can start the roll forward. You cannot start Tivoli Storage Productivity Center at this point, because the database does not open until you complete some type of roll forward.

Roll forward options in DB2 can be complex. We do not intend to provide a complete guide to DB2 roll forward recovery. We describe how to roll forward in two ways:

Roll forward to the end of the logs

This way rolls forward from the restore point through all available log files to the most recent consistent point-in-time. If you are using an old backup and there are many logs through which to roll, this method can take time.

Roll forward to a point-in-time

With a point-in-time roll forward, you can specify a specific point-in-time for the roll forward process to stop, complete, and allow the database to open.

Rolling the database forward to the end of the logs

To roll the database forward to the end of all of the logs after a restoration, type the following commands in the DB2 command-line processor as shown in Figure 16-30:

rollforward database TPCDB to end of logs and complete

When each command completes, it returns an audit of the process.

Tip: The last committed transaction time is displayed in a UTC-0 time format even if your local time zone is, for example, PDT (UTC-8).

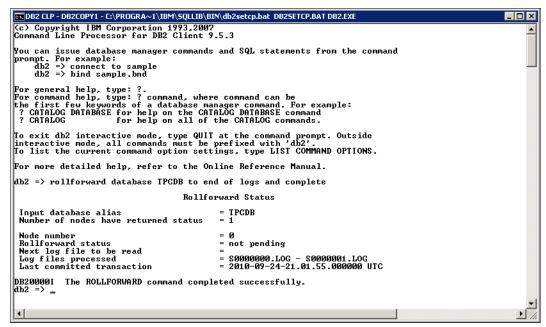


Figure 16-30 Roll forward TPCDB to the end of the logs and complete

When complete, proceed to "Restarting the Tivoli Storage Productivity Center services" on page 684.

Rolling the database forward to a point-in-time

Tip: By default, DB2 uses UTC-0 time for the point-in-time roll forward. Add the **use local time** flag to the command if you want to specify a time in your local time zone.

Use the following steps to roll the database forward to a given point-in-time after the restore:

- 1. Use the DB2 command-line processor as shown in Figure 16-31 to enter the **rollforward** command. In this example, we rolled the TPCDB database forward to a few minutes after the restore time. We entered the time using the **use local time** option.
- 2. Enter the point-in-time as YYYY-MM-DD-HH.MM.SS format.

The command for the TPCDB database is as follows:

rollforward database TPCDB to 2010-09-26-14.52 using local time and complete

```
os. DB2 CLP - DB2COPY1 - C:\PROGRA~1\IBM\SQLLIB\BIN\db2setcp.bat DB2SETCP.BAT DB2.EXI
                                                                                                                                                                            _ 🗆 🗆 ×
(c) Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.5.3
You can issue database manager commands and SQL statements from the command
    ompt. For example:
db2 => connect to sample
db2 => bind sample.bnd
For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
? CATALOG DATABASE for help on the CATALOG DATABASE command
? CATALOG for help on all of the CATALOG commands.
  o exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.
For more detailed help, refer to the Online Reference Manual.
db2 => rollforward database TPCDB to 2010-09-26-14.52 using local time and complete
                                                          Rollforward Status
                                                                      = TPCDB
= 1
 Input database alias
Number of nodes have returned status
 Node number
Rollforward status
Next log file to be read
Log files processed
Last committed transaction
                                                                        = not pending
                                                                      =
= $0000000.LOG - $0000114.LOG
= 2010-09-26-14.51.03.000000 Local
DB20000I The ROLLFORWARD command completed successfully. db2 => \_
```

Figure 16-31 Roll forward the TPCDB to point-in-time and complete

Notice that the actual *last committed transaction* time differs slightly from the time that is requested in the roll forward. This time is the closest that DB2 can get to the requested time and still keep the database in a consistent state.

Restarting the Tivoli Storage Productivity Center services

After you complete the restore operation and the roll forward for the Tivoli Storage Productivity Center database, restart the Tivoli Storage Productivity Center service to bring the server back online. To do this on a WindowsTivoli Storage Productivity Center server, issue the commands shown in Example 16-24 in a command window.

Example 16-24 Windows commands to start Tivoli Storage Productivity Center

```
net start "IBM Tivoli Storage Productivity Center - Data Server" net start "IBM WebSphere Application Server V6.1 - DeviceServer"
```

16.12.3 Potential agent issues after the restore process

Following the restore process of the Tivoli Storage Productivity Center database to a previous state, there is always a possibility that new agents were deployed to machines after the time of the restore. If this happens, certain agents, which are not registered in the Tivoli Storage Productivity Center server and which are now running with a older version of the database, are out in the environment. To correct this problem, instruct the orphaned agents to re-register themselves with the Tivoli Storage Productivity Center server.

Reinstall the agents with the **force** parameter by using the **Agent** command or a deployment job from the GUI.

16.13 Backup scheduling and automation

This section provides information about the method to use to back up the database and how often to do it.

16.13.1 Frequency of full TPCDB backups

How often to take a full backup of your Tivoli Storage Productivity Center database depends on how critical the TPCDB data is to your business.

Run the full TPCDB backup once a week. If it is significantly critical, implement a TPCDB backup strategy to accommodate your business needs. For example, a full TPCDB backup can be scheduled every weekend, and incremental backups (not explained in this chapter) can be scheduled every week day.

For details about this subject, see *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441.

16.13.2 TPCDB backup automation

Taking backups of the TPCDB database can be automated. Several available options for this task are as follows:

- Windows Task Scheduler
- ▶ DB2 Administration Server's scheduler
- ► Tivoli Storage Manager Backup-Archive Scheduler.

See the following publications for detailed information about this subject:

- ► IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference, SC27-2441
- ► Backing Up DB2 Using IBM Tivoli Storage Management, SG24-6247



Tivoli Storage Productivity Center database considerations

This chapter describes Tivoli Storage Productivity Center database repository growth and how you consolidate and eventually prune the amount of history stored in the Tivoli Storage Productivity Center database repository.

This chapter also covers basic tips for improving DB2 performance, which can be beneficial for organizations that are planning large scale deployments of Tivoli Storage Productivity Center.

This chapter covers the following topics:

- Database tuning
- ► Repository database backup and recovery
- ► Repository size planning for:
 - Disk subsystem performance collection
 - SAN Fabric performance collection
 - Statistical data: Tivoli Storage Productivity Center for Data
- History aggregation
- Preparing for database Backups

17.1 Database tuning

The Tivoli Storage Productivity Center installer configures certain DB2 parameters at installation time. However, for many environments, these initial settings were found to be inadequate.

Table 17-1 shows the settings that we suggest.

Table 17-1 DB2 recommended settings for Tivoli Storage Productivity Center

Parameter description	Default value	DB2 parameter name	Tivoli Storage Productivity Center recommendation
Default application heap size	10240	applheapsz	20480
Database heap size	1000	dbheap	1500
Log buffer size	8	logbufsz	512
Log file size	2500	logfilsiz	20000
Number of primary log files	8	logprimary	150
Number of secondary log files	100	logsecond	100
Maximum DB files open per application	64	maxfilop	1024
mon_heap_sz: includes space that cannot be used for storage volumes or virtual storage volumes.	132	mon_heap_sz	1024
Statement heap size	10240	stmtheap	20480
IBMDEFAULTBP: Buffer pool size	250	IBMDEFAULTBP	4000
TPCBFPDATA: Buffer pool size	250	TPCBFPDATA	25000
TPCBFPKEYS: Buffer pool size	250	TPCBFPKEYS	2500
TPCBFPTEMP: Buffer pool size	250	TPCBFPTEMP	1000

17.1.1 Setting DB2 variables

To set the suggested value, complete the following steps:

- 1. Do one of the following steps, depending on the operating system:
 - On Windows:

Launch a DB2 command-line processor window into your DB2 server. See Figure 17-1.

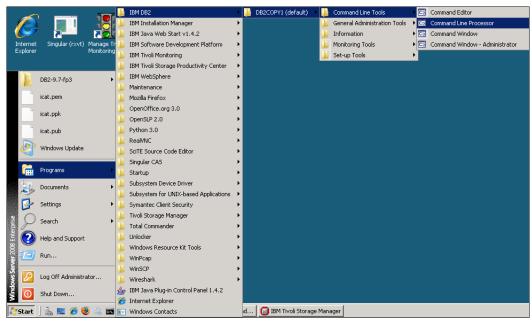


Figure 17-1 Launch a DB2 command line processor

A command-line processor opens, as shown in Figure 17-2.

Figure 17-2 DB2 command line processor

- On AIX or Linux:
 - i. Log in to your DB2 server and switch to the DB instance owner ID (usually db2inst1) or source the instance profile:
 - . /home/db2inst1/sqllib/db2profile
 - ii. Initiate the DB2 command-line processor:

db2

2. Connect to the Tivoli Storage Productivity Center database (TPCDB):

```
connect to tpcdb
```

3. Tune the TPCDB database:

```
update db cfg for tpcdb using APPLHEAPSZ 20480 update db cfg for tpcdb using DBHEAP 1500 update db cfg for tpcdb using LOGFILSIZ 20000 update db cfg for tpcdb using LOGPRIMARY 150 update db cfg for tpcdb using LOGSECOND 100 update db cfg for tpcdb using LOGBUFSZ 512 update db cfg for tpcdb using MAXFILOP 1024 update db cfg for tpcdb using STMTHEAP 20480
```

4. Tune the buffer pool:

```
alter bufferpool IBMDEFAULTBP immediate size 4000 alter bufferpool TPCBFPDATA immediate size 25000 alter bufferpool TPCBFPKEYS immediate size 2500 alter bufferpool TPCBFPTEMP immediate size 1000
```

5. Close the database connection:

connect reset

17.1.2 Tuning the database manager

This section explains how to tune the database manager.

Updating the database manager settings

Complete the following steps:

- 1. Attach to the database instance:
 - On Windows:

attach to DB2

On AIX or Linux:

attach to db2inst1

2. Update DB2 database manager settings:

```
update dbm cfg using MON HEAP SZ 1024
```

3. Close the connection:

detach

4. The new settings go into effect the next time that the database closes and opens. Stop Tivoli Storage Productivity Center and restart it to use the new settings.

Another method to use the new settings is to reboot the Tivoli Storage Productivity Center server.

Restarting Tivoli Storage Productivity Center services

Restart the Tivoli Storage Productivity Center services by issuing the commands in a command window as shown in one of the following examples:

► On Windows, use Example 17-1.

Example 17-1 Windows commands to restart Tivoli Storage Productivity Center

```
net stop "IBM Tivoli Storage Productivity Center - Data Server"
net stop "IBM WebSphere Application Server V6.1 - DeviceServer"
net start "IBM Tivoli Storage Productivity Center - Data Server"
net start "IBM WebSphere Application Server V6.1 - DeviceServer"
```

► On AIX, use Example 17-2.

Example 17-2 AIX commands to restart Tivoli Storage Productivity Center

```
stopsrc -s TSRMsrv1
/<usr or opt>/IBM/TPC/device/bin/aix/stopTPCF.sh
startsrc -s TSRMsrv1
/<usr or opt>/IBM/TPC/device/bin/aix/startTPCF.sh
```

► On Linux, use Example 17-3.

Example 17-3 Linux commands to restart Tivoli Storage Productivity Center

```
/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 stop
/<usr or opt>/IBM/TPC/device/bin/linux/stopTPCF.sh
/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 start
/<usr or opt>/IBM/TPC/device/bin/linux/startTPCF.sh
```

17.1.3 Changing DB2 active logs directory

By default, the Tivoli Storage Productivity Center database and its DB2 active logs are stored on the same file system. You can achieve performance improvements by placing the logs on a separate file system or a separate disk drive to balance the I/O requirements of both tasks.

To move the logs for the TPCDB database to a new location, use the following steps:

1. Choose a new log path location (for this example, we use the following locations):

On Windows: D:\DB2_active_logs\TPCDBOn AIX or Linux: /var/DB2/active_logs/TPCDB

2. Start a DB2 command-line processor (Figure 17-3).

Figure 17-3 DB2 command line processor

- 3. Issue the following command, depending on the operating system:
 - On Windows:

update db cfg for TPCDB using newlogpath D:\DB2 active logs\TPCDB

On AIX or Linux:

update db cfg for TPCDB using newlogpath /var/DB2/active_logs/TPCDB

4. The new log path goes into effect the next time that the database closes and opens. Stop Tivoli Storage Productivity Center and restart it to use the new log path. Another method to use the new log path is to reboot the Tivoli Storage Productivity Center server.

Stopping Tivoli Storage Productivity Center services

Stop the Tivoli Storage Productivity Center services by issuing the commands in a command window as shown in one of the following examples:

► On Windows, use Example 17-4.

Example 17-4 Windows commands to stop Tivoli Storage Productivity Center

net stop "IBM Tivoli Storage Productivity Center - Data Server" net stop "IBM WebSphere Application Server V6.1 - DeviceServer"

On AIX, use Example 17-5.

Example 17-5 AIX commands to stop Tivoli Storage Productivity Center

stopsrc -s TSRMsrv1
/<usr or opt>/IBM/TPC/device/bin/aix/stopTPCF.sh

➤ On Linux, use Example 17-6.

Example 17-6 Linux commands to stop Tivoli Storage Productivity Center

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 stop
/<usr or opt>/IBM/TPC/device/bin/linux/stopTPCF.sh

Starting the Tivoli Storage Productivity Center services

Start the Tivoli Storage Productivity Center server by issuing the commands in a command window, as shown in one of the following examples:

► On Windows, use Example 17-7.

Example 17-7 Windows commands to start Tivoli Storage Productivity Center

net start "IBM Tivoli Storage Productivity Center - Data Server" net start "IBM WebSphere Application Server V6.1 - DeviceServer"

► On AIX, use Example 17-8.

Example 17-8 AIX commands to start Tivoli Storage Productivity Center

startsrc -s TSRMsrv1
/<usr or opt>/IBM/TPC/device/bin/aix/startTPCF.sh

▶ On Linux, use Example 17-9.

Example 17-9 Linux commands to start Tivoli Storage Productivity Center

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 start
/<usr or opt>/IBM/TPC/device/bin/linux/startTPCF.sh

17.2 Repository database sizing

This section is a guide to estimating the storage requirements for the Tivoli Storage Productivity Center repository database. This section is not designed to be an exact tool. Instead, it focuses on the primary data types that consume the majority of space within the Tivoli Storage Productivity Center repository database.

17.2.1 Introduction

You can divide the sizing information into three sections:

- Storage subsystem performance collection data:
 - IBM DSxxxx, XIV, and non-IBM subsystems
 - IBM SAN Volume Controller (SVC) and Storwize V7000 systems
- ► SAN Fabric performance collection data
- ► Tivoli Storage Productivity Center for Data analysis collection records

Tivoli Storage Productivity Center collects performance data from each device at timed intervals. For instance, you might decide to collect performance information from your production DS8000 at 5-minute intervals. This interval consumes much more repository storage than if you choose to sample your production DS8000 at 30-minute intervals.

Storage subsystems SVC, Storwize V7000, XIV, DS3000, DS4000, DS6000, and ESS all support the performance monitoring sample interval for performance data collection tasks from a minimum interval of five minutes to a maximum interval of 60 minutes.

Important: If you plan to collect performance data from a large number of devices at small time intervals, the infrastructure that you establish must be able to support it. The CIMOMs that collect the data will need to work harder, and the Tivoli Storage Productivity Center repository database will need to support a higher number of transactions to keep pace.

You might be required to perform additional tuning of the DB2 database as monitored devices increase or if you plan to use a high sample rate.

As performance data arrives at Tivoli Storage Productivity Center, it is stored in the Tivoli Storage Productivity Center repository database. Over time, Tivoli Storage Productivity Center creates hourly and daily averages of this data. The averaged data requires less storage space in the repository over a longer period of time. It also makes reporting over a longer time period more meaningful and easier to display.

For this scenario, you plot the performance of a subsystem for the last 30 days. If you do so using 15-minute samples, 2,800 plot points would exist, resulting in a graph that is both jagged and difficult to plot.

Because Tivoli Storage Productivity Center averages the 15-minute samples into both hourly and daily samples, known as *aggregates*, you can choose to plot the same 30-day period with only 30 or 720 points on the graph, resulting in a less jagged and more legible graph.

The process of collecting the individual samples into hourly and daily samples is called *history aggregation*, and Tivoli Storage Productivity Center has a configuration panel that controls how much history you keep over time.

Important: The history aggregation process is a global setting, which means that the values set for history retention are applied to *all* performance data from *all* devices. You cannot set history retention on an individual device basis.

Figure 17-4 shows the Tivoli Storage Productivity Center panel for setting the history retention for performance monitors and other types of collected statistics.

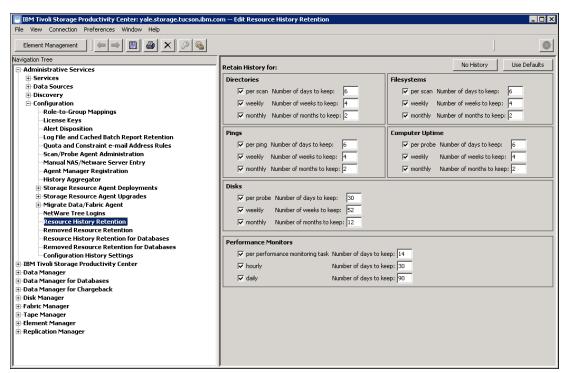


Figure 17-4 Setting the history retention for performance monitors

The Performance Monitors values in Figure 17-4 are as follows:

► Per performance monitoring task:

The value that you set here defines the number of days that Tivoli Storage Productivity Center keeps individual data samples for all of the devices that send performance data. The example shows 14 days. When per sample data reaches this age, Tivoli Storage Productivity Center permanently deletes it from the database.

Increasing this value allows you to look back at device performance at the most granular level at the expense of consuming more storage space in the Tivoli Storage Productivity Center repository database.

Data held at this level is good for plotting performance over a small time period but not for plotting data over many days or weeks because of the number of data points. Consider keeping more data in the hourly and daily sections for longer time period reports.

This check box determines whether history retention is on or off:

- If you clear or do not select the check box, the data collected by the data collection schedule is retained permanently. This might cause a large amount of data to accumulate in the database repository over time.
- If you select the check box and enter the value 0 for items under Performance Monitors, the data already collected by those schedules is deleted and no more data is collected.

► Hourly:

This value defines the number of days that Tivoli Storage Productivity Center holds performance data that has been grouped into hourly averages. Hourly average data potentially consumes less space in the database. For example, if you collect performance

data from an SVC at 15-minute intervals, retaining the hourly averages requires four times less space in the database.

This check box determines whether history retention is on or off. Select the check box and enter 0 to not keep the historical data.

▶ Daily:

This value defines the number of days that Tivoli Storage Productivity Center holds performance data that has been grouped into daily averages. After the defined number of days, Tivoli Storage Productivity Center permanently deletes records of the daily history from the repository.

Daily averaged data requires 24 times less space in the data for storage compared to hourly data. This is at the expense of granularity; however, plotting performance over a longer period (perhaps weeks or months) becomes more meaningful.

This check box determines whether history retention is on or off. Select the check box and enter 0 to not keep the historical data.

17.2.2 Storage subsystem performance data sizing

A significant difference exists in the sizing calculations between the SVC or Storwize V7000 and other subsystems, both IBM and non-IBM. For this reason, the sizing tables are separated in this section.

Sizing the repository for XIV, DSxxxx, and non-IBM subsystems

You can use the example worksheet in Table 17-2 to gain an understanding of the likely storage requirements that are needed for the repository to hold a given amount of data.

Table 17-2 shows working examples for four storage subsystems in an environment and the amount of storage space that performance collection uses for each example. The total figure represents the amount of storage needed for the "per sample" data. Continue through this section to calculate the complete amount of storage needed for hourly and daily history types.

The calculation method example for XIV_Production is as follows:

 $60/5 \times 24 = 288$ samples per day x 500 volumes x 200 bytes per sample = 28,800,000 bytes

Table 17-2 Per sample repository database sizing for XIV and subsystems (not IBM)

(a) Subsystem name	(b) Number of volumes (LUNs) sampled	(c) Performance collection interval (minutes)	(d) Performance data record size	(e) Daily amount of data collected (60/(c) x 24) x (b) x (d) = (e)
XIV_production	500	5	200 bytes	28800000
EMC_remote	320	5	200 bytes	14,400,000
		(f) Total required pe	er day	43200000
		(g) Number of days to retain per sample = 14 days		907 MB
		(f) x (g)/1,024,000 + 50%		

Tip: The final value includes an additional 50%, which provides for DB2 table indexes and other database overhead.

As you can see, the amount of space that is required increases dramatically as the sample rate increases. Remember this concept when you plan the appropriate sample rate for your environment.

Next, use Table 17-3 to calculate the amount of storage that is needed to hold the performance data for the hourly and daily history averages. When complete, add together the totals from Table 17-2 on page 696 and Table 17-3 to give you the total repository requirement for these types of storage subsystems as seen in Table 17-4.

The calculation method example for XIV_production is as follows:

500 volumes x **200** bytes per sample x 24 = 2,400,000 bytes for hourly history average

Table 17-3 Hourly and daily repository database sizing for XIV and storage (not IBM)

(a) Subsystem name	(b) Number of volumes sampled (LUNs)	(c) Performance data record size (bytes)	(d) Hourly requirement	(e) Daily requirement
<u></u>			(b) x (c) x 24	(b) x (c)
XIV_production	500	200	2,400,000	100,000
EMC_remote	320	200	1,536,000	64,000
		Daily totals	3,936,000	164,000
		Hourly retention days = 30	(f) 118,080,000	
		Daily retention days = 90		(g) 14760000
			Total MB (f) + (g)/ 1,024,000 + 50%	190 MB

Table 17-4 shows the total Tivoli Storage Productivity Center repository space required for XIV and non-IBM storage subsystems. The total Tivoli Storage Productivity Center repository space is the sum of the totals of both Table 17-2 on page 696 and Table 17-3.

Table 17-4 Total repository space required

Total space required MB	
907 + 190	
1,097 MB	

Sizing the repository for SVC and Storwize V7000 performance collection

Tivoli Storage Productivity Center collects a large number of performance metrics from SVC and Storwize V7000 devices to take advantage of many performance metrics reported by the device software.

Complete Table 17-5 for each SVC or Storwize V7000 that you will monitor in the environment. The table assumes a fixed sample rate is chosen for all SVCs and Storwize V7000s. If you plan to monitor some at 15-minute intervals and others at 30-minute intervals, you must fill out this table twice (one table for each chosen sample rate), and then, add the two tables together to give you an overall total.

Table 17-5 Repository sizing for SVC and Storwize V7000

Subsystem	Number of VDisks	Number of MDisks	I/O groups	MDisk groups	Cluster pairs
TEC_SVC	900	500	1	4	1
SVC_DR	3,000	1,500	2	6	2
Totals	3,900	2,000	3	10	3
Record size (bytes)	198	78	500	128	492
Byte totals	772,200	156,000	1,500	1,280	1,476
			(a) Sample ra	ate (bytes)	932,456
			(b) Hourly an sample rate (60/15) x	nount @ 15-minute	3,729,824
			(c) Daily amo	ount	89,515,776
			(d) 30-day re (b) x 24 x	tention of samples 30	2,685,473,280
			(e) 30-day re 24 x (a) x	tention of hourly 30	671,368,320
			(f) 90-day ret (a) x 90	ention of daily	83,921,040
				tal required (MB) /1,024,000 + 50%	4,999 MB

Important: Notice that the overall value in (g) adds 50% to the amounts calculated through the table. The majority of this overhead takes the DB2 table indexes for this data plus database page overhead into account.

Sizing the repository for SAN fabric performance data

This section describes sizing for SAN fabric performance collection. Fabric port record sizes per data sample are relatively large. We also observed that the indexing requirements for this data also tend to be high.

We base the byte sizing that we provide here on real observations of actual database utilization over a period of time.

Table 17-6 shows all switches that are sampled at 5-minute intervals. If you plan to monitor some switches at one rate and other switches at another rate, create a separate table for each rate. The final value includes a 50% uplift for indexing and DB2 storage overhead.

Table 17-6 SAN switch performance repository data sizing

Switch name	(a) Number of ports	(b) Size (bytes)	(c) Sample rate (minutes)	(d) Hourly amount (bytes) (60/(c)) x (a) x (b)
TEC_Switch_1	32	400	5	153,600
TEC_Switch_2	32	400	5	153,600
Remote_Switch_1	64	400	5	307,200
Remote_switch_2	64	400	5	307,200
Totals	(e) 192	(f) Total sample rate per hour		921,600
		(g) 30 days retain sample rate (f) x 24 x 30		663,552,000
		(h) 30 days retain hourly rate (e) x (b) x 24 x 30		55,296,000
		(i) 90 days retain daily rate (e) x (b) x 90		6,912,000
		Overall Total MB (g) + (h) + (i)/1,024,000 + 50%		1038.2 MB

Sizing the repository for Tivoli Storage Productivity Center for Data requirements

Repository sizing for Tivoli Storage Productivity Center for Data is more difficult to accurately model because of the dynamic nature of the collected data. Performance data collection sizing is simple in that it collects a set amount of data at regular intervals.

However, with Tivoli Storage Productivity Center for Data, a policy or profile collects a variable amount of data from each monitored server based on what and how much data of a matching type is found on each machine.

Key factors in sizing Tivoli Storage Productivity Center for Data are as follows:

- ► Total number of operating system registered users storing files
- ► Total number of file systems monitored
- ► Total number of various file types (that is, *.txt, *.exe, *.doc, *.mp3, and so on)
- Number of machines with data agents deployed and collecting data
- ► Total number of file names collected and stored for reporting

Key largest repository tables:

- ▶ T_STAT_USER_HIST User history file
- ► T_STAT_FTYPE_HIST File type history
- ► T_STAT_FILE Stored file names

Figure 17-5 on page 702, Table 17-8 on page 701, and Table 17-9 on page 703 help to estimate the worst-case sizing for these key tables. See Table 17-7 for detail on estimating the user history for the Tivoli Storage Productivity Center for Data repository.

Table 17-7 Estimating the user history repository requirement

Statistic name	Number of file systems covered	Number of users covered	Days to keep scan history	Number of weeks of scan history	Number of months of scan history
Custom_stat	300	800	30	52	24
UNIX_stat	250	1500	30	52	24
Windows_stat	500	2000	30	52	24
Totals	(a) 1050	(b) 4300	(c) 90	(d) 156	(e) 72
			Worst case to requirement a x b x (c + d +	(bytes)	64,609,650,000
			Realistic exp		6,460,965,000
			Divide 1,024	,000 = MB	6,310 MB

Tip: Unlike the performance tables, we must estimate much more here. For example, there might be 500 file systems covered by the Windows_stat and 2000 users with data across the 500 file systems, but not all 500 file systems have files owned by the 2000 users. Likely only a subset of file systems have data for all 2000 users. For this reason, the realistic figure is reduced to only 10% of the worst case figure. You might want to change the 10% factor to your specific requirements.

Use Table 17-8 to calculate the repository space that is required to store file type history information.

Estimating the file type history buildup is more accurate than estimating the user table history, because the data entering this table is more constant for a given profile.

Table 17-8 Estimating the file type repository requirement

Statistic profile name	Number of file types	Number of TPC agents covered	Days to keep scan history	Number of weeks of scan history	Number of months of scan history
Win_types	50	200	30	52	24
UNIX_servers	50	150	60	52	24
Media files	30	50	60	52	24
Totals	(a) 130	(b) 400	(c) 150	(d) 156	(e) 72
			Total - bytes a x b x (c + d +	- e) x 55 bytes	1,081,080,000
			Total MB - to	tal/1,024,000	1,056 MB

The third Tivoli Storage Productivity Center for Data repository table of significant size is the T_STAT_FILE table. This table holds a record of the file names, which have been collected by profiles for largest, most obsolete, orphan files, and so forth.

Tip: If you plan to use Tivoli Storage Productivity Center for Data for duplicate file spotting or to archive specific files for you, you will likely increase the number of file names that each agent will collect.

When completing Table 17-9, the "Total file names per agent" will be the total of all types, shown in Figure 17-5. In this example, it is 1,800 file names per agent.

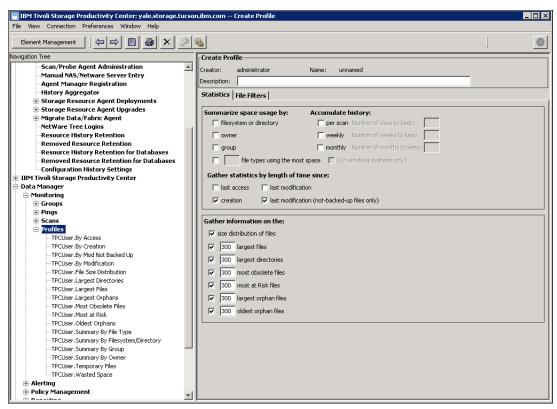


Figure 17-5 Adding up all file names

Table 17-9 shows the estimating of the file name repository requirements.

Table 17-9 Estimating the file name repository requirement

Statistic profile name	(a) Total file names collected per agent	(b) Number of agents to which this profile applies	Total files per statistic a x b
Duplicate file spot	2,000	500	1,000,000
Control audio files	200	150	30,000
Archive old data	200	50	10,000
		Total files in table	1,040,000
		Size (bytes) = Total x 250 bytes	420,000,000
		Size/1,024,000 = MB	410 MB

The final step for sizing the Tivoli Storage Productivity Center for Data repository is to total the three tables and add an overhead for the default statistics. The average overhead for the default statistic types is provided at 1.5 MB per Tivoli Storage Productivity Center agent. Therefore the formula is as follows:

Default TPC for Data overhead = Total agents x 1.5 MB

An example is as follows:

 $1,000 \times 1.5 = 1,500 \text{ MB}$

Enter this value in Table 17-10.

Table 17-10 Tivoli Storage Productivity Center for Data repository total

Source	Amount in MB
User history	6,310
File type history	1,056
File names	410
Default statistics overhead	1,500
Total requirement (MB)	9,276

17.3 Repository calculation templates

This section has blank versions of the worksheets that are used in 17.2, "Repository database sizing" on page 693 to calculate the DB2 repository space requirements for a given environment. Use these worksheets to help you size your individual requirements.

17.3.1 Worksheet: Sizing SVC and Storwize V7000 performance collection

Use Table 17-11. See "Sizing the repository for SVC and Storwize V7000 performance collection" on page 697 for a working example.

Table 17-11 Repository sizing for SVC and Storwize V7000

Subsystem	Number of VDisks	Number of MDisks	I/O groups	MDisk groups	Cluster pairs
Totals					
Record size (bytes)	198	78	500	128	492
Byte totals					
			(a) Sample ra	ate (bytes)	
			(b) 15-minute (60/15) x		
			(c) Daily amo	ount	
			(d) 14-day re (b) x 24 x	tention of samples	
			(e) 30-day re 24 x (a) x	tention of hourly 30	
			(f) 90 days of (a) x 90		
				tal required (MB) /1,024,000 + 50%	

17.3.2 Worksheet: Sizing performance collection for XIV, DSxxxx, and non-IBM subsystems

See "Sizing the repository for XIV, DSxxxx, and non-IBM subsystems" on page 696 for a working example of this table.

Table 17-12 is the first of two worksheets needed to calculate the repository space that is required for these types of subsystems.

Table 17-12 Per sample repository database sizing for XIV, DSxxxx, and non-IBM subsystems

(a) Subsystem name	(b) Number of volumes (LUNs) sampled	(c) Performance collection interval (minutes)	(d) Performance data record size	(e) Daily amount of data collected (60/(c) x 24) x (b) x (d) = (e)
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
			200	
	•	(f) Total required pe	er day	
		(g) Number of days sample = 14 days (f) x (g)/1,024,000		

Table 17-13 is the second table needed to calculate repository space required for these types of subsystems.

Table 17-13 Hourly and daily repository database sizing for XIV, DSxxxx, and non-IBM storage

(a) Subsystem name	(b) Number of volumes sampled (LUNs)	(c) Performance data record size (bytes)	(d) Hourly requirement (b) x (c) x 24	(e) Daily requirement (b) x (c)
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		200		
		Daily totals		
		Hourly days x 30	(f)	
		Daily days = 90		(g)
			Total MB (f) + (g)/1,024,000 + 50%	

17.3.3 Worksheet: Sizing SAN switch performance collection

See "Sizing the repository for SAN fabric performance data" on page 698 for a working example of how to calculate the amount of storage that is required to hold SAN switch performance data. See Table 17-14.

If you monitor SAN switches at various time intervals, use a separate worksheet for each sample rate chosen.

Table 17-14 SAN switch performance repository data sizing

Switch name	(a) Number of ports	(b) Size (bytes)	(c) Sample rate (minutes)	(d) Hourly amount (bytes) (60/(c)) x (a) x (b)
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
		400		
Totals	(e)	(f) Total sample	rate per hour	
		(g) 14 days reta (f) x 24 x 14		
		(h) 30 days reta (e) x (b) x 24		
		(i) 90 days retain (e) x (b) x 90		
		Overall Total ME (g) + (h) + (i)/1,0		

17.3.4 Worksheet: Sizing Tivoli Storage Productivity Center for Data repository

See 17.2, "Repository database sizing" on page 693 for working examples of these tables. See Table 17-15 for estimating user history.

Table 17-15 Estimating the user history repository requirement

Statistic name	Number of file systems covered	Number of users covered	Days to keep scan history	Number of weeks of scan history	Number of months of scan history
Totals	(a)	(b)	(c)	(d)	(e)
			Worst case to requirement a x b x (c + d +	(bytes)	
			Realistic expectation - reduce to 10% of worst		
			Divide 1,024,	.000 = MB	

Table 17-16 is used for estimating file type repository requirement.

Table 17-16 Estimating the file type repository requirement

Statistic profile name	Number of file types	Number of TPC agents covered	Days to keep scan history	Number of weeks of scan history	Number of months of scan history
Totals	(a)	(b)	(c)	(d)	(e)
			Total - bytes a x b x (c + d +	+ e) x 55 bytes	
			Total MB - to	tal/1,024,000	

Table 17-17 is used for estimating file name repository requirement.

Table 17-17 Estimating the file name repository requirement

Statistic profile name	(a) Total file names collected per agent	(b) Number of agents to which this profile applies	Total files per statistic a x b
		Total files in table	
		Size (bytes) = Total x 250 bytes	
		Size/1,024,000 = MB	

The default overhead for the default statistic types is 1.5 MB per Tivoli Storage Productivity Center for Data agent. Therefore, consider the following formula:

Default TPC for Data overhead = Total agents x = 1.5 MB

An example is as follows:

 $1,000 \times 1.5 = 1,500 \text{ MB}$

Enter this figure in Table 17-18.

Table 17-18 Tivoli Storage Productivity Center for Data repository total

Source	Amount in MB
User history	
File type history	
File names	
Default statistics overhead	
Total requirement MB	

17.4 Preparing for database backups

This section discusses the necessary preparation for database backups.

17.4.1 Scripts provided

This book provides the scripts listed in Table 17-19 as-is for your convenience. They are simple scripts to give you an understanding of how the process of backup works for file system type backups and Tivoli Storage Manager backups. Use these scripts as a basis for your own processes, and modify them as necessary.

Table 17-19 Example scripts provided

Script name	Function
TPC_backup_offline_file.bat TPC_backup_offline_file (Unix)	Backs up the Tivoli Storage Productivity Center database offline to a file system
TPC_backup_offline_tsm.bat TPC_backup_offline_tsm (Unix)	Backs up the Tivoli Storage Productivity Center database offline to Tivoli Storage Manager
TPC_backup_online_file.bat TPC_backup_online_file (Unix)	Backs up the Tivoli Storage Productivity Center database online to a file system
TPC_backup_online_tsm.bat TPC_backup_online_tsm (Unix)	Backs up the Tivoli Storage Productivity Center database online to Tivoli Storage Manager
database_list_offline_file.txt	Lists the databases to back up for an offline file system
database_list_offline_tsm.txt	Lists the databases to back up for offline to Tivoli Storage Manager
database_list_online_file.txt	Lists the databases to back up for online file system
database_list_online_tsm.txt	Lists the databases to back up for online to Tivoli Storage Manager

17.4.2 Database backup

This section describes the high-level information that you need to understand about DB2 UDB before you can plan a backup strategy for securing Tivoli Storage Productivity Center V4.

Backup types

The two primary methods of backing up DB2 databases are offline and online:

- ► Offline backup (sometimes known as cold backup) is when all database access is terminated, and the database is closed. The backup then runs stand-alone before the database is restarted and access-enabled. This is the simplest type of backup to set up, configure, and maintain.
- Online backup (sometimes known as hot backup) is when all user and application database access continues to run while the backup process takes place. This type of backup provides for continuous availability of the database and the applications that require it. This is a more complex type of backup to set up, configure, and maintain.

Directing the backup output destination

You can direct the database backup output to a number of destinations from within DB2:

File system:

Direct output to normal file system structure flat files. Then, you can copy these files to removable tape for added security or back them up with products, such as Tivoli Storage Manager or other widely available similar tools.

Tape:

Send output directly to tape if the tape device is directly attached to the server hardware.

► Tivoli Storage Manager:

DB2 sends output directly to Tivoli Storage Manager through direct integration between the two products. If a Tivoli Storage Manager environment exists within your organization, you can back up directly to it by installing the Tivoli Storage Manager Backup/Archive client and client API on the same computer that hosts the Tivoli Storage Productivity Center DB2 database.

► XBSA:

DB2 directs output to an application that is compliant with X/Open Backup Services Application (X/BSA), such as Legato NetWorker.

▶ Vendor DLL:

DB2 directs output to a third-party vendor-written interface API.

Attention: This chapter focuses on the file system and Tivoli Storage Manager backup destinations.

Database logging

DB2 UDB uses log files to keep a sequential record of all database changes. The log files are specific to DB2 UDB activity. The logs record the database activity in transactions. If a crash occurs, you use logs to play back or redo committed transactions during recovery.

The two types of logging are as follows:

Circular logging (default)

This type is the simplest and is the default logging type that Tivoli Storage Productivity Center uses.

Archive logging

This type of logging enables online backup and roll-forward recovery of a database to a point-in-time. It is, however, more complex to manage.

17.4.3 Database backup method considerations

This section considers the merits of offline backup methods compared to online backup methods for the Tivoli Storage Productivity Center database. The default method of backup for Tivoli Storage Productivity Center is to use offline backup.

Offline backup advantages and disadvantages

Here we discuss the relative advantages and disadvantages of this method:

- Advantages of offline backup are as follows:
 - Simple: You can perform offline backup with DB2 logging set to the default circular method.
 - DB2 skills: Offline backup requires a minimum amount of DB2 skills DB2 to perform, because it is the simplest method of backup.
 - Logging: Circular logs are the simplest to manage and maintain.
- ► Disadvantages of offline backup are as follows:
 - Stopped Tivoli Storage Productivity Center server services: The offline method involves stopping all Tivoli Storage Productivity Center server services on a regular basis (typically daily) to perform the backup. This regular outage might not be acceptable to all organizations that want to use Tivoli Storage Productivity Center.
 - Missed performance data collection: If you have set up Tivoli Storage Productivity Center to continuously collect disk subsystem and SAN fabric performance statistics, you lose data points for the duration that Tivoli Storage Productivity Center is down each day for backup. You can minimize the impact of this loss by scheduling the backup at a time when the monitored equipment statistics are of little importance from a reporting perspective. This loss of data points might not be acceptable to all organizations wanting to use Tivoli Storage Productivity Center.
 - Missed events: Tivoli Storage Productivity Center monitors the infrastructure and alerts you about events, such as failures within a SAN fabric. You risk missing critical events if the events occur when you stop the Tivoli Storage Productivity Center server services for the backup process.

Online backup advantages and disadvantages

Here we discuss the relative advantages and disadvantages of this method:

- Advantages of online backup are as follows:
 - Greater availability: You do not need to stop and start the Tivoli Storage Productivity
 Center server services on a daily basis for the backup operation. Online backups do
 not interrupt user access to the database while the backup operation is in progress.
 - No missed events: Tivoli Storage Productivity Center monitors the infrastructure and alerts you about events, such as failures within a SAN fabric. Using online backup ensures that Tivoli Storage Productivity Center is able to respond quickly to critical events at any time of the day.
 - Uninterrupted performance collection: You experience no interruption or missing data points in the collection of performance data from disk subsystems and SAN fabrics.
- ► Disadvantages of online backup are as follows:
 - More DB2 skills required: Archive logging is a more advanced method of DB2 operation, and administering archive logging requires more skills.

Tivoli Storage Productivity Center software update process can fail: Our testing found that Tivoli Storage Productivity Center software updates that alert the database layout can fail. You must revert to circular logging to perform updates, and then switch back to archive logging.



Tivoli Storage Productivity Center database backup on Linux

This chapter explains how to plan for backing up and restoring the Tivoli Storage Productivity Center database that resides in DB2 on the Tivoli Storage Productivity Center server in a Linux environment.

The chapter covers both offline backup (cold backup) and online backup (hot backup) of the database along with the merits of each type.

The Tivoli Storage Productivity Center product does not provide any extra backup and recovery tools in addition to those tools already provided with the DB2 product. This chapter is not intended to be a comprehensive guide to all functions of backup and recovery built into DB2. For details about this subject, see the *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441.

18.1 Before you start

We suggest that you read Chapter 17, "Tivoli Storage Productivity Center database considerations" on page 687 before you continue with this chapter. To improve the performance of your DB2 instance, tune your database settings by following the instructions provided in this section, and review 17.4.3, "Database backup method considerations" on page 711.

You can also estimate the storage requirements for the Tivoli Storage Productivity Center repository database, and plan the sizing of your backup volumes accordingly.

18.2 Common backup setup steps

Complete the following setup steps for both file system and Tivoli Storage Manager backups:

 Configure the DB2 history file to retain the number of backup versions that you want to retain. Your organization might already have a policy for how many versions you need to keep.

Change the DB2 num_db_backups parameter and set the value to the number of backup versions that you require. You must also set the rec_his_retentn parameter to a value of -1. By setting this value to -1, the rec_his_retentn parameter follows the value set in the num_db_backups parameter.

Important: For this change to take effect, Tivoli Storage Productivity Center service must be stopped and started. This restarting does not necessarily need to happen directly after you change the parameter.

- 2. Log into your DB2 server and switch to the DB instance owner ID (usually db2inst1) or source the instance profile:
 - . /home/db2inst1/sqllib/db2profile

Then, initiate the DB2 command-line processor:

db2

3. Example 18-1 shows how to set the num_db_backups value to 4 versions and rec_his_retentn to -1 for both the Tivoli Storage Productivity Center database.

Issue the commands at the db2 => prompt in the command-line processor window.

Example 18-1 DB2 commands to configure how many backup versions to keep

```
connect to TPCDB
update db cfg using num_db_backups 4
update db cfg using rec_his_retentn -1
disconnect TPCDB
quit
```

Important: When you set new values for num_db_backups and rec_his_retentn, the new values are not effective until you stop all database connections.

 Restart Tivoli Storage Productivity Center to make the changes effective. You can either reboot the server, or alternatively stop and start the services, as shown in Example 18-2.

Example 18-2 Linux commands to stop and start Tivoli Storage Productivity Center services

```
/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 stop
/<usr or opt>/IBM/TPC/device/bin/linux/stopTPCF.sh

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 start
/<usr or opt>/IBM/TPC/device/bin/linux/startTPCF.sh
```

18.3 Offline backup to file system setup steps

This section describes how to set up offline backup for the Tivoli Storage Productivity Center server database to flat files in a file system. Because the offline backup method is the default method for Tivoli Storage Productivity Center, there is little DB2 configuration needed before you can perform a backup.

Important: Ensure that you perform the steps in 18.2, "Common backup setup steps" on page 714 and these steps.

The steps are as follows:

Choose a location to use for the DB2 backup output. Choose a directory that has enough
free space to hold the number of backups that you plan to retain. It is best to use a
separate file system rather than the file system that contains the DB2 database.

You can choose to use a location that is a remotely mounted CIFS or NFS, so that the backup data is secured to another server, perhaps at another location in your organization.

This example uses /var/TPC_database_backups.

Important: DB2 does not create this directory for you. Create this directory before you attempt a backup, and make sure that user db2inst1 has write permissions.

2. Create a batch script to control the backup process.

Two files are used:

- /root/TPCBKP/TPC backup offline file

This script file (shown in Example 18-3) runs the backup.

Example 18-3 File /root/TPCBKP/TPC_backup_offline_file

```
#!/bin/bash

#This is a sample backup script
#To backup TPC offline
#To disk filesystems
. /home/db2inst1/sqllib/db2profile

echo "Stopping Tivoli Storage Productivity Center services"
echo "-----"
echo
```

```
/opt/IBM/TPC/data/server/tpcdsrv1 stop
/opt/IBM/TPC/device/bin/linux/stopTPCF.sh
echo
echo "Starting backup of the DB2 database"
echo "-----"
db2 force application all
db2 $(cat /root/TPCBKP/database list offline file.txt )
echo "Restarting Tivoli Storage Productivity Center services"
echo "-----"
echo
/opt/IBM/TPC/data/server/tpcdsrv1 start
/opt/IBM/TPC/device/bin/linux/startTPCF.sh
echo
echo
echo "Offline backup process complete"
echo "-----"
exit 0
```

```
Remember: Make the script executable by using the chmod command: chmod +x /root/TPCBKP/TPC_backup_offline_file
```

/root/TPCBKP/database_list_offline_file.txt

This file (shown in Example 18-4) is DB2 scripted list of databases to back up.

Example 18-4 File /root/TPCBKP/database_list_offline_file.txt

backup database TPCDB to /var/TPC database backups without prompting

See 18.7.1, "Performing an offline backup to a file system" on page 727 to run an offline backup.

18.4 Offline backup to Tivoli Storage Manager setup steps

This section describes the steps necessary to set up an offline backup of the Tivoli Storage Productivity Center server database to a Tivoli Storage Manager server. The backup to Tivoli Storage Manager is a little more complex to set up but does not require you to set aside large amounts of local disk space for backup versions on the Tivoli Storage Productivity Center server.

In this section, we assume the following requirements:

- You have a basic working knowledge of Tivoli Storage Manager.
- ► An operational Tivoli Storage Manager server already exists to which you can send backup data.
- Your Tivoli Storage Manager administrator has defined storage, which will receive the backups, to the policies.
- You have already installed a Tivoli Storage Manager Backup-Archive client on the Tivoli Storage Productivity Center server, and you have configured it to do standard file backups.
- You have installed the Tivoli Storage Manager API Client on the Tivoli Storage Productivity Center server.
- You used default installation paths for Tivoli Storage Manager.

Stop and restart: You must stop Tivoli Storage Productivity Center and DB2 as part of this configuration process. Plan this exercise at a time when you can restart Tivoli Storage Productivity Center.

Use the following steps, described in this section, to configure DB2 to Tivoli Storage Manager integration:

- 1. "Adding new variables to Linux" on page 717
- 2. "Configuring Tivoli Storage Manager option file and password" on page 717
- 3. "Restarting DB2" on page 719
- 4. "Creating an offline backup to Tivoli Storage Manager script" on page 719

18.4.1 Adding new variables to Linux

Table 18-1 shows a list of Tivoli Storage Manager API environment variables to add to Linux. The listed values assume a default installation of Tivoli Storage Manager on the Tivoli Storage Productivity Center server.

Table 18-1 System environment variables

Environment variable name	Value
DSMI_DIR	/usr/tivoli/tsm/client/api/bin64
DSMI_CONFIG	/home/db2inst1/tsm/dsm.opt
DSMI_LOG	/home/db2inst1/tsm

Set the environment variables for the API client. As shown in Example 18-5, add the following lines in the \$HOME/.profile file for the DB2 instance administrator (usually /home/db2inst1/.profile):

Example 18-5 Add Tivoli Storage Manager variables to the DB2 profile

```
echo 'export DSMI_DIR=/opt/tivoli/tsm/client/api/bin64
export DSMI_CONFIG=/home/db2inst1/tsm/dsm.opt
export DSMI_LOG=/home/db2inst1/tsm
' >> /home/db2inst1/sqllib/db2profile
```

Important:

- ► If you are using a 32-bit version of DB2, use /opt/tivoli/tsm/client/api/bin instead of /opt/tivoli/tsm/client/api/bin64.
- If it does not exist, create /home/db2inst1/tsm directory, and make sure that user db2inst1 is the owner.

18.4.2 Configuring Tivoli Storage Manager option file and password

This section describes the steps necessary to configure the Tivoli Storage Manager option file dsm.opt and then set the Tivoli Storage Manager password so that the DB2 backup process can communicate with the Tivoli Storage Manager API.

Important: At this stage, your Tivoli Storage Manager client must be already registered with a Tivoli Storage Manager server. If the Tivoli Storage Manager server is accepting open registrations, just by starting the Tivoli Storage Manager client GUI or command line you will be asked for a password to register your client. If the Tivoli Storage Manager server is using closed registration, you will need the Tivoli Storage Manager administrator to register your client.

The steps are as follows:

1. Edit the dsm.sys file, which is in the following location:

/opt/tivoli/tsm/client/api/bin64/

32-bit version: If you are using a 32-bit version of DB2, edit the following file instead: /usr/tivoli/tsm/client/api/bin/dsm.sys

Make sure that the client option PASSWORDACCESS is set to GENERATE as shown in Figure 18-1.

```
[root@tpcblade6-11 bin64]# cat /opt/tivoli/tsm/client/api/bin64/dsm.sys
```

Servername TSMsrv1 COMMMethod TCPIP

TCPPort 1500

TCPSERVERADDRESS tsmsrv1.storage.tucson.ibm.com

PASSWORDACCESS GENERATE

Figure 18-1 Contents of the dsm.sys file

2. Create or edit the dsm.opt file, in the /home/db2inst1/tsm/ location.

The dsm.opt file only needs to have one line in it which is a reference to the server stanza in the dsm.sys file, which in our case is TSMsrv, as shown in Figure 18-2.

Figure 18-2 Contents of the dsm.opt file

- 3. Set the Tivoli Storage Manager password so that DB2 can authenticate with the Tivoli Storage Manager server when DB2 performs a backup or restore operation:
 - a. Run the dsmapipw command as shown in Figure 18-3 on page 719.
 - b. Enter the current and new Tivoli Storage Manager password. You can reuse the existing Tivoli Storage Manager password.

Important: You must run the **dsmapipw** command even if you do not intend to change the Tivoli Storage Manager password. Running this command registers the password with the Tivoli Storage Manager API. Registering this password in the setup phase means that a DB2 operator can perform backup and restore operations without needing to know the Tivoli Storage Manager client password. If a Tivoli Storage Manager administrator changes or resets the Tivoli Storage Manager password, you need to run the **dsmapipw** command again.

Figure 18-3 Running the dsmapipw command

Important: Check that files dsierror.log and dsm.opt in the /home/db2inst1/tsm directory are owned by the DB2 instance owner (db2inst1) to avoid errors during the backup process.

18.4.3 Restarting DB2

Now that you have completed the configuration steps, restart Tivoli Storage Productivity Center and DB2 as shown in Example 18-6 to ensure that the environment variables are picked up.

Example 18-6 Linux commands to restart Tivoli Storage Productivity Center services and DB2

```
. /home/db2inst1/sqllib/db2profile

/opt/IBM/TPC/data/server/tpcdsrv1 stop
/opt/IBM/TPC/device/bin/linux/stopTPCF.sh

db2 force applications all
db2stop
db2start

/opt/IBM/TPC/data/server/tpcdsrv1 start
/opt/IBM/TPC/device/bin/linux/startTPCF.sh
```

18.4.4 Creating an offline backup to Tivoli Storage Manager script

Create two files:

► The first file is the script that you run (Example 18-7):

/root/TPCBKP/TPC_backup_offline_tsm

Example 18-7 File /root/TPCBKP/TPC_backup_offline_tsm

```
#!/bin/bash

#This is a sample backup script
#To backup TPC offline
#To disk filesystems
```

. /home/db2inst1/sqllib/db2profile

```
echo "Stopping Tivoli Storage Productivity Center services"
echo "------"
echo
/opt/IBM/TPC/data/server/tpcdsrv1 stop
/opt/IBM/TPC/device/bin/linux/stopTPCF.sh
echo
echo "Starting backup of the DB2 database"
echo "-----"
db2 force application all
db2 $(cat /root/TPCBKP/database_list_offline_tsm.txt )
echo
echo "Restarting Tivoli Storage Productivity Center services"
echo "------"
echo
/opt/IBM/TPC/data/server/tpcdsrv1 start
/opt/IBM/TPC/device/bin/linux/startTPCF.sh
echo
echo "Offline backup process complete"
echo "------"
exit 0
```

```
Remember: Make the script executable by using the chmod command:

chmod +x /root/TPCBKP/TPC_backup_offline_tsm
```

► The second file (Example 18-8) is the DB2 scripted list of databases to back up: root/TPCBKP/database list offline tsm.txt

```
Example 18-8 File /root/TPCBKP/database_list_offline_tsm.txt
```

backup database TPCDB use tsm without prompting

18.5 Online backup to Tivoli Storage Manager setup steps

This section describes the steps for configuring the Tivoli Storage Productivity Center database to enable for online backup to Tivoli Storage Manager. The significant difference between online and offline backup is the need to enable archive logging on the database. As we describe in "Database backup method considerations" on page 711, operating in this mode provides many backup and recovery benefits at the expense of increased complexity in the database operation.

Important: You must stop Tivoli Storage Productivity Center to perform these tasks:

- ▶ DB2 requires a full backup of each database before you can start the Tivoli Storage Productivity Center database again after these reconfiguration steps. We include the instructions to perform a full backup of the database. Allow time in your outage planning for the backup to complete.
- ► Also, complete the steps in 18.2, "Common backup setup steps" on page 714 to set the number of backup versions that you want to retain in the history file.

Be sure you consider the advantages and disadvantages of archive logging before you continue with this setup. For full details of DB2 logging methods, see the DB2 product manuals. Also see *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441, for detailed information about this subject.

Considerations:

- ▶ If you set up DB2 for online backup to Tivoli Storage Manager, you cannot easily change to an online backup to file system. You need to choose between these methods, because you are setting the destination for the archive logging process. If you decide in the future to change to the online file system method, you will need to reconfigure DB2 to send the archive logs to file system. To complete the task, this reconfiguration requires restarting Tivoli Storage Productivity Center.
- ▶ It is possible to perform an online backup to file system and have the archive logs going to Tivoli Storage Manager. However, this approach is not desirable because the difficulty of managing and tracking information makes this a poor practice.

Set up and test DB2 to Tivoli Storage Manager integration before you attempt the steps in this section. Use 18.4, "Offline backup to Tivoli Storage Manager setup steps" on page 716. When you are satisfied that DB2 is communicating with Tivoli Storage Manager and you have performed at least one successful offline backup, return to this section.

18.5.1 DB2 parameter changes for archive logging to Tivoli Storage Manager

To set up archive logging to Tivoli Storage Manager, complete the following tasks:

 You must make a number of parameter choices for the configuration of archive logging as shown in Table 18-2. These parameters determine where DB2 keeps its log files. Be sure that the db2inst1 user is the owner of all log directories.

Table 18-2 DB2 parameters

DB2 parameter	Example value	Comment
Primary log path	/var/DB2/active_logs	This location is where DB2 keeps the current logs for the database. For best performance, place these logs on a separate volume than the volume that holds the data.
Failed log path	/var/DB2/failed_logs	This location is where DB2 put log files if the archive process fails. This can happen if Tivoli Storage Manager is down or unreachable when DB2 tries to send a log file to Tivoli Storage Manager.

2. Stop Tivoli Storage Productivity Center by using the commands in Example 18-9.

Example 18-9 Linux commands to stop Tivoli Storage Productivity Center

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 stop
/<usr or opt>/IBM/TPC/device/bin/linux/stopTPCF.sh

- 3. Log in to your DB2 server and switch to the DB instance owner ID (usually db2inst1) or source the instance profile:
 - . /home/db2inst1/sqllib/db2profile

Then, initiate the DB2 command line processor:

db2

4. Issue the commands from Example 18-10 at the command-line processor window. Substitute your chosen values for the parameters that form part of the UPDATE DB CFG command. See Table 18-2. The final command performs an offline backup of the database.

Important: The database backup is required after this reconfiguration, and the DB2 database will not open again until the database backup is completed.

Example 18-10 DB2 command to configure archive logging to Tivoli Storage Manager

CONNECT TO TPCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR TPCDB USING logarchmeth1 TSM failarchpath
/var/DB2/failed_logs newlogpath /var/DB2/active_logs

BACKUP DATABASE TPCDB USE TSM

QUIT

Attention: Verify that the following directories exist and are owned by the db2inst1 user:

/var/DB2/logs/DB2_failed_logs
/var/DB2/logs/DB2 active logs

5. When the database backup is complete, you can restart Tivoli Storage Productivity Center. Issue the commands shown in Example 18-11.

Example 18-11 Start Tivoli Storage Productivity Center

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 start
/<usr or opt>/IBM/TPC/device/bin/linux/startTPCF.sh

18.5.2 Creating an online backup script for Tivoli Storage Manager

Create two files:

► The first file is the script (Example 18-12) that you run to start the backup:

```
/root/TPCBKP/TPC_backup_online_tsm
```

Example 18-12 File /root/TPCBKP/TPC_backup_online_tsm

```
#!/bin/bash

#This is a sample backup script
#To backup TPC online
#To Tivoli Storage Manager

. /home/db2inst1/sqllib/db2profile
echo
echo "Starting backup of the DB2 database"
echo "-----"
db2 $(cat /root/TPCBKP/database_list_online_tsm.txt)
echo
echo "Online backup process complete"
echo "------"
exit 0
```

```
Remember: Make the script executable by using the chmod command:

chmod +x /root/TPCBKP/TPC backup online tsm
```

The second file is the DB2 scripted list of databases (Example 18-13) to back up:

```
/root/TPCBKP/database_list_online_tsm.txt
```

Example 18-13 File /root/TPCBKP/database_list_online_tsm.txt

backup database TPCDB online use tsm without prompting

18.6 Online backup to a file system setup steps

Performing online backups to a file system requires you to set up archive logging to a file system also. When operating with this method, DB2 does not clean up old and no longer necessary archive log files. Therefore, you need to put processes in place to clean up old log files after a specific amount of time to prevent the system from filling up. You also need to plan for this amount of space. Over a number of weeks, the log space that is required for a Tivoli Storage Productivity Center database can become many times larger than the database.

To be able to restore an online DB2 database taken two weeks ago, for example, you must have log files going back to that same date that you can restore. An online DB2 database backup is not stand-alone, because you cannot restore the online DB2 database backup without at least some logs for it to roll forward to a consistent state.

Important:

- ► Although switching between a backup destination of online to a file system and online to Tivoli Storage Manager is a straightforward process, switching the logging path is not as easy. To switch the logging from Tivoli Storage Manager to a file system requires a stop and a start of the database and, therefore, a restart of the Tivoli Storage Productivity Center services.
- Choose either a Tivoli Storage Manager backup or a file system backup, and stay with that specific method.

18.6.1 Setting up DB2 archive logging to a file system

Set up DB2 archive logging to a file system by using the following steps:

 To configure archive logging, you must make several parameter choices. The parameters in Table 18-3 determine where DB2 will keep its log files, how many log files to keep, and the size of the log files. Be sure that the db2inst1 user is the owner of all log directories.

Table 18-3 DB2 parameters for archive logging to a file system

DB2 parameter	Example value	Comment
Primary log path	/var/DB2/active_logs	The location where DB2 will keep the current logs for the database. For best performance, place the logs on a separate volume than the data.
Archive log path	/var/DB2/archive_logs/TPCDB	The location where DB2 will archive log files for the TPCDB database.
Failed log path	/var/DB2/failed_logs	The location where DB2 will put log files if the archive process fails, which can happen if the file system for the primary logs fills up. Choose a location that is <i>not</i> on the same file system as the archive logs.

2. Choose a file system path (Table 18-4) to store the DB2 database backups. Ensure that that the directory is owned by user db2inst1.

Table 18-4 File system location for database backups

Database backup path
/var/TPC_database_backups

3. Stop Tivoli Storage Productivity Center by using the commands in Example 18-14.

Example 18-14 Linux commands to stop Tivoli Storage Productivity Center

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 stop
/<usr or opt>/IBM/TPC/device/bin/linux/stopTPCF.sh

- 4. Log in to your DB2 server and switch to the DB instance owner ID (usually db2inst1) or source the instance profile:
 - . /home/db2inst1/sqllib/db2profile

Next, initiate the DB2 command-line processor:

db2

Issue the commands shown in Example 18-15 in the command-line processor window.
 Substitute your values for the parameters that form part of the UPDATE DB CFG command. See Table 18-3. Note that the final command performs an offline backup of the database.

Important: The offline backup of the database is required after the reconfiguration; the DB2 database will not open until the backup is complete.

Example 18-15 DB2 command to configure archive logging to a file system

CONNECT TO TPCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR TPCDB USING logarchmeth1 DISK:/var/DB2/archive_logs/TPCDB failarchpath /var/DB2/failed_logs newlogpath /var/DB2/active_logs

BACKUP DATABASE TPCDB TO /var/TPC_database_backups

6. When both database backups complete, you can restart Tivoli Storage Productivity Center. Issue the commands shown in Example 18-16.

Example 18-16 Start Tivoli Storage Productivity Center

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 start
/<usr or opt>/IBM/TPC/device/bin/linux/startTPCF.sh

18.6.2 Creating an online backup script to a file system

Create two files:

► The first file is the script (Example 18-17) that you run to start the backup:

```
/root/TPCBKP/TPC_backup_online_file
```

Example 18-17 File /root/TPCBKP/TPC_backup_online_file

```
#!/bin/bash

#This is a sample backup script

#To backup TPC online

#To disk filesystems

. /home/db2inst1/sqllib/db2profile
echo
echo "Starting backup of the DB2 database"
echo "------"
db2 $(cat /root/TPCBKP/database_list_online_file.txt)
echo
echo "Online backup process complete"
echo "------"
exit 0
```

```
Remember: Make the script executable by using the chmod command: chmod +x /root/TPCBKP/TPC_backup_online_file
```

The second file is the DB2 scripted list of databases to back up (Example 18-18):

```
/root/TPCBKP/database_list_online_file.txt
```

```
Example 18-18 File /home/root/TPCBKP/database_list_online_file.txt
```

backup database TPCDB online to /var/TPC database backups without prompting

18.7 Performing offline database backups

This section describes how to perform the offline backup of the Tivoli Storage Productivity Center database.

Running an offline DB2 database backup takes Tivoli Storage Productivity Center out of service for the period of the backup. This impacts data collections for your devices, and you might miss other infrastructure events.

Make sure that you understand the impact of stopping Tivoli Storage Productivity Center in your environment before proceeding. If your environment cannot tolerate regularly stopping Tivoli Storage Productivity Center for a backup operation, consider configuring and using online backup.

18.7.1 Performing an offline backup to a file system

Important: You must complete the initial steps as detailed in 18.2, "Common backup setup steps" on page 714 before you can start to perform offline backups.

To perform an offline backup to a file system, run the **TPC_backup_offline_file** script (Figure 18-4). This script stops the Tivoli Storage Productivity Center process, performs a backup of the DB2 database to the location specified in the scripts, and then restarts the services.

```
[root@tpcblade6-11 ~]# /root/TPCBKP/TPC backup offline file
Stopping Tivoli Storage Productivity Center services
Setting Variables for SANM
Stopping server1 with default options
ADMU0116I: Tool information is being logged in file
/opt/IBM/TPC/device/apps/was/profiles/deviceServer/logs/server1/stopServer.log
ADMU0128I: Starting tool with the deviceServer profile
ADMU3100I: Reading configuration for server: server1
ADMU3201I: Server stop request issued. Waiting for stop status.
ADMU4000I: Server server1 stop completed.
Starting backup of the DB2 database
DB20000I The FORCE APPLICATION command completed successfully.
DB21024I This command is asynchronous and may not be effective immediately.
Backup successful. The timestamp for this backup image is : 20100930174149
Restarting Tivoli Storage Productivity Center services
_____
Setting Variables for SANM
Starting server1 for Device Manager
9/30/10 5:41:57 PM GEN0198I: Server starting
ADMU0116I: Tool information is being logged in file
/opt/IBM/TPC/device/apps/was/profiles/deviceServer/logs/server1/startServer.log
ADMU0128I: Starting tool with the deviceServer profile
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 28700
Offline backup process complete
```

Figure 18-4 Running an offline backup to a file system

18.7.2 Performing an offline backup to Tivoli Storage Manager

Important: You must complete the initial setup steps described in 18.4, "Offline backup to Tivoli Storage Manager setup steps" on page 716 before you can start offline backups.

Running an offline DB2 database backup takes Tivoli Storage Productivity Center out of service for the period of the backup. Make sure it is acceptable to take Tivoli Storage Productivity Center out of service before you proceed.

To perform an offline backup to Tivoli Storage Manager, run the TPC_backup_offline_tsm script as shown in Figure 18-5.

```
[root@tpcblade6-11 ~]# /root/TPCBKP/TPC backup offline tsm
Stopping Tivoli Storage Productivity Center services
Setting Variables for SANM
Stopping server1 with default options
ADMU0116I: Tool information is being logged in file
/opt/IBM/TPC/device/apps/was/profiles/deviceServer/logs/server1/stopServer.log
ADMU0128I: Starting tool with the deviceServer profile
ADMU3100I: Reading configuration for server: server1
ADMU3201I: Server stop request issued. Waiting for stop status.
ADMU4000I: Server server1 stop completed.
Starting backup of the DB2 database
-----
DB20000I The FORCE APPLICATION command completed successfully.
DB21024I This command is asynchronous and may not be effective immediately.
Backup successful. The timestamp for this backup image is: 20100930174744
Restarting Tivoli Storage Productivity Center services
Setting Variables for SANM
Starting server1 for Device Manager
9/30/10 5:48:47 PM GEN0198I: Server starting
ADMU0116I: Tool information is being logged in file
/opt/IBM/TPC/device/apps/was/profiles/deviceServer/logs/server1/startServer.log
ADMU0128I: Starting tool with the deviceServer profile
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 30481
Offline backup process complete
-----
```

Figure 18-5 Running an offline backup to Tivoli Storage Manager

18.8 Performing online database backup

This section describes how to run online backups of the Tivoli Storage Productivity Center database. By running the backups online (or hot), it is not necessary to stop the Tivoli Storage Productivity Center service, which means that the backup process does not interrupt processes, such as performance data collection.

Important: Before you can perform online DB2 backup, you must configure Tivoli Storage Manager to DB2 communication and DB2 archive logging to Tivoli Storage Manager or DB2 to file system logging.

If you plan to use online backup to Tivoli Storage Manager, perform the setup steps in 18.4, "Offline backup to Tivoli Storage Manager setup steps" on page 716 and 18.5, "Online backup to Tivoli Storage Manager setup steps" on page 720.

18.8.1 Performing an online database backup to Tivoli Storage Manager

This section describes how to run an online backup of the Tivoli Storage Productivity Center database to Tivoli Storage Manager. You must have already run the appropriate setup steps.

To do a backup, run the TPC_backup_online_tsm script as shown in Figure 18-6.

```
[root@tpcblade6-11 ~]# /root/TPCBKP/TPC_backup_online_tsm

Starting backup of the DB2 database
------

Backup successful. The timestamp for this backup image is : 20100930175157

Online backup process complete
-------
```

Figure 18-6 Running an online backup to Tivoli Storage Manager

18.8.2 Performing an online backup to a file system

This section describes how to run an online backup of the Tivoli Storage Productivity Center database to file system output files. You must already have completed the setup steps described in 18.6, "Online backup to a file system setup steps" on page 723.

To do the backup, run the TPC_backup_online_file script as seen in Figure 18-7.

```
[root@tpcblade6-11 ~]# /root/TPCBKP/TPC_backup_online_file

Starting backup of the DB2 database
-----

Backup successful. The timestamp for this backup image is : 20100930175729

Online backup process complete
------
```

Figure 18-7 Running an online backup to file system

18.9 Other backup considerations

Apart from the DB2 database, a number of important files and directories must be backed up up to preserve the state of a Tivoli Storage Productivity Center server installation.

Back up all files under the Tivoli Storage Productivity Center installation directory. In addition, the Tivoli GUID is stored in /etc/Tivoli/TSRM location. Be sure that you also preserve this Tivoli GUID.

The important Tivoli Storage Productivity Center server directories to secure are as follows:

- ► TPC_Server_install_dir/config/
- ► TPC Server install dir/data/config/
- ► TPC Server install dir/device/conf/

These directories contain the various configuration files for your installation. Saving these directories is important, because they might be customized configurations and not the default configurations.

18.10 Managing database backup versions

This section is an overview of the maintenance processes for which you need to plan. DB2 does not prune older versions automatically. Therefore, you must maintain the number of Tivoli Storage Productivity Center database backup versions that you want on a file system or Tivoli Storage Manager.

18.10.1 Managing backup versions for a file system

This section describes what you need to know to manage DB2 backups that were performed to disk.

Deleting unwanted backups

DB2 does not automatically manage the deletion of the unwanted database backups or archive logs from the file system. You need to create a maintenance plan to delete the old backups. If you plan to use online backup to a file system, you need to create a plan to delete the old archive logs.

How DB2 organizes backups on a file system

When you do a backup to a file system, you supply the backup script with a path to use. Since the release of Version 9.1, DB2 creates backup files with a file name containing a datestamp that details when the backup was taken.

Figure 18-8 shows the backup directory of the TPCDB database. The directory contains backups taken at separate times on 30 September 2010. DB2 timestamps all backups in this way; every time a backup is made, a new file is created:

► The first part of the file name starts as follows:

```
TPCDB.O.DB2.NODE0000.CATN0000
```

► The last part of the file name consists of the date in yyyyMMDD format, as in the following example:

20100930

Plan to delete old backup files to suit the requirements of your backup and recovery policy.

```
[root@tpcblade6-11 ~]# ls /var/TPC_database_backups
TPCDB.0.db2inst1.NODE0000.CATN0000.20100930141335.001
TPCDB.0.db2inst1.NODE0000.CATN0000.20100930174149.001
TPCDB.0.db2inst1.NODE0000.CATN0000.20100930173433.001
TPCDB.0.db2inst1.NODE0000.CATN0000.20100930175729.001
TPCDB.0.db2inst1.NODE0000.CATN0000.20100930173932.001
```

Figure 18-8 DB2 backup directory

18.10.2 Managing archive log files on a file system

If you plan to do online backups, you must configure DB2 to use archive logging. If you plan to do online backups to disk, you must also maintain the archive logs directory on a regular basis.

Figure 18-9 shows the directory structure for logging on to the TPCDB database. Over time, this directory fills up with logs. If your recovery policy is to keep backup versions for five days, you must keep logs in this directory for at least the same period of time. The reason is, to make the recovery valid, you cannot restore an online backup without logs from the same date and time.

Notice that the directory that holds the logs is named C0000000, which is the log cycle number. If you restore the database, the cycle number increments by one and starts in C0000001 and so on. Ensure that any automated deletion process that you implement can handle this numbering.

```
[root@tpcblade6-11 ~]# ls -R /var/DB2/archive_logs/TPCDB/

/var/DB2/archive_logs/TPCDB/:
db2inst1

/var/DB2/archive_logs/TPCDB/db2inst1:
TPCDB

/var/DB2/archive_logs/TPCDB/db2inst1/TPCDB:
NODE0000

/var/DB2/archive_logs/TPCDB/db2inst1/TPCDB/NODE0000:
C00000000

/var/DB2/archive_logs/TPCDB/db2inst1/TPCDB/NODE0000/C00000000:
S0000001.LOG S0000002.LOG S0000003.LOG S0000004.LOG S0000005.LOG S0000006.LOG
```

Figure 18-9 DB2 archive logs

18.10.3 Managing backup versions that you store in Tivoli Storage Manager

This section describes how to maintain, view, and delete backup data and archive logs that you have sent to Tivoli Storage Manager. DB2 does not automatically prune backup versions and log files from Tivoli Storage Manager. You need to use the db2adut1 tool to perform these housekeeping functions.

Tip: This section is not intended to be a comprehensive guide to the **db2adut1** tool. The intent here is to detail the commands that you likely will use on a regular basis to maintain the data that is held in Tivoli Storage Manager.

Purpose of the db2adutl command

The command line tool, db2adut1, communicates with Tivoli Storage Manager through its API interface. Use this tool to interrogate the backup and archive log data that is stored in Tivoli Storage Manager at any one time, verify that you no longer require old backups, and delete unnecessary old backups.

Reasons to use db2adutl

When DB2 stores a backup session in Tivoli Storage Manager, DB2 always stores the backup session with a unique file name, which is the timestamp when the backup was made. This means that these backup versions never get superseded by a new version with the same file name. The backup files remain "active" versions in Tivoli Storage Manager, and, therefore, Tivoli Storage Manager never deletes the backup versions.

Use the command, db2adut1, to select unwanted backup versions and tell Tivoli Storage Manager to flag them as "inactive." In this way, Tivoli Storage Manager then deletes them over time based on the standard policy rules that the Tivoli Storage Manager administrator set.

You handle DB2 archive logs differently. They are stored in Tivoli Storage Manager as "archive" data, which means Tivoli Storage Manager retains them for a set period of time based on its policies. Although you can use **db2adut1** to explicitly remove DB2 archive logs, if Tivoli Storage Manager archive retention policy is set appropriately, this is not necessary.

Important: Make sure that the Tivoli Storage Manager archive retention policy that you use to store the DB2 logs is set for a sufficient period of time to allow recovery of your oldest database backup. However, you also want to make sure that the policy for the retention period is not so long that it wastes storage space in Tivoli Storage Manager.

How to guery backups held in Tivoli Storage Manager

Next, we explain how to guery backups that are held in Tivoli Storage Manager.

Important: The **db2adut1** command is normally in the bin directory of the DB2 instance owner, which is usually the following location:

/home/db2inst1/sqllib/bin/db2adutl.

Before using the command, log in with the db2inst1 user ID.

The db2adut1 command to guery database backup versions are as follows:

► The following command lists all the database versions and the logs that are held for all databases stored in Tivoli Storage Manager and TPCDB, in this case:

db2adut1 query

► The following command lists all database versions and logs for the TPCDB database. Note that the database name is case-sensitive and is in capital letters:

db2adutl query database TPCDB

Figure 18-10 shows the sample output from this command. The output shows that two database backups are stored in Tivoli Storage Manager and six archive logs.

► The following command has a shorter output. It lists only the database backup versions and the archive logs:

db2adutl query full

```
[db2inst1@tpcblade6-11 ~]$ db2adutl query database TPCDB
Query for database TPCDB
Retrieving FULL DATABASE BACKUP information.
   1 Time: 20100930204416 Oldest log: S0000025.LOG DB Partition Number: 0
Sessions: 2
   2 Time: 20100930204152 Oldest log: S0000024.LOG DB Partition Number: 0
Sessions: 2
   3 Time: 20100930203906 Oldest log: S0000024.LOG DB Partition Number: 0
Sessions: 1
   4 Time: 20100930202923 Oldest log: S0000023.LOG DB Partition Number: 0
Sessions: 1
   5 Time: 20100930202350 Oldest log: S0000022.LOG DB Partition Number: 0
Sessions: 1
   6 Time: 20100930201854 Oldest log: S0000021.LOG DB Partition Number: 0
Sessions: 1
   7 Time: 20100930200626 Oldest log: S0000020.LOG DB Partition Number: 0
Sessions: 1
   8 Time: 20100930194948 Oldest log: S0000018.LOG DB Partition Number: 0
Sessions: 2
   9 Time: 20100930193637 Oldest log: S0000017.LOG DB Partition Number: 0
Sessions: 2
  10 Time: 20100930192744 Oldest log: S0000017.LOG DB Partition Number: 0
Sessions: 1
  11 Time: 20100930191237 Oldest log: S0000013.LOG DB Partition Number: 0
Sessions: 2
  12 Time: 20100930184747 Oldest log: S0000008.LOG DB Partition Number: 0
Sessions: 2
-- 8< ---- OUTPUT CLIPPED -- 8< ----
Retrieving LOG ARCHIVE information.
  Log file: S0000000.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-09-30-15.57.39
  Log file: S0000002.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-09-30-18.22.12
  Log file: S0000024.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-09-30-20.32.56
  Log file: S0000025.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-09-30-20.35.33
```

Figure 18-10 Sample output from a db2adutl query database TPCDB command

Deleting backup versions held in Tivoli Storage Manager

Important: Because Tivoli Storage Manager does not allow the root user to delete backups that are created by the db2inst1 instance, log in with user ID db2inst1 before trying this commands.

The following commands and examples show how to delete database backup versions that are held in Tivoli Storage Manager:

► The following command deletes backup versions from Tivoli Storage Manager that are older than ninety days. This type of command is useful, because you can easily script it to run each day to remove the older backup.

```
db2adutl delete full older than 90 days
You can also specify a database name:
db2adutl delete full older than 90 days database TPCDB
```

Figure 18-11 gives you an example of running this command.

► The following command deletes all backup versions from Tivoli Storage Manager, except for the last five versions. Again, this command is useful when scripting an automatic process.

```
db2adutl delete full keep 5
You can also specify a database name:
db2adutl delete full keep 5 database TPCDB
```

```
[db2inst1@tpcblade6-11 ~]$ db2adutl delete full keep 5 database TPCDB
Query for database TPCDB
Taken at: 20100930201854 DB Partition Number: 0
                                                  Sessions: 1
                                                  Sessions: 1
 Taken at: 20100930200626 DB Partition Number: 0
 Taken at: 20100930194948 DB Partition Number: 0
                                                    Sessions: 2
 Taken at: 20100930193637 DB Partition Number: 0 Sessions: 2
 Taken at: 20100930192744 DB Partition Number: 0 Sessions: 1
 Taken at: 20100930191237 DB Partition Number: 0 Sessions: 2
 Taken at: 20100930184747 DB Partition Number: 0
                                                    Sessions: 2
    Do you want to delete these backup images (Y/N)? Y
     Are you sure (Y/N)? Are you sure (Y/N)? Y
The current delete transaction failed. You do not have
sufficient authorization. Attempting to deactivate
backup image(s) instead...
Success.
Retrieving INCREMENTAL DATABASE BACKUP information.
 No INCREMENTAL DATABASE BACKUP images found for TPCDB
Retrieving DELTA DATABASE BACKUP information.
 No DELTA DATABASE BACKUP images found for TPCDB
```

Figure 18-11 Example of a db2adutl delete full keep 5 database TPCDB command

Managing DB2 archive log files in Tivoli Storage Manager

The commands described in this section are examples of how to delete database archive logs from Tivoli Storage Manager.

Important:

- ▶ Be careful when you delete archive log files. If you delete logs that are still needed for some of your backup versions, you render those backups *useless*.
- Archive logs only exist in Tivoli Storage Manager if you have configured archive logging so that online backup is possible.
- Ask the Tivoli Storage Manager administrator to configure Tivoli Storage Manager to delete the archive logs on a regular basis by configuring the Tivoli Storage Manager "archive copy group" that DB2 uses. Set a retention period that suits your needs. If you use a general purpose archive copy group, Tivoli Storage Manager might keep all archive logs for several years causing unnecessary usage of the storage in your Tivoli Storage Manager environment.

To delete archive logs, first query the Tivoli Storage Manager server to establish which logs you want to delete.

To query the Tivoli Storage Manager server for the TPCDB database, issue the following command (the output is shown in Figure 18-12):

db2adutl query database TPCDB

```
[db2inst1@tpcblade6-11 ~]$ db2adutl query database TPCDB
Query for database TPCDB
Retrieving FULL DATABASE BACKUP information.
    1 Time: 20100930204416    Oldest log: S0000025.LOG    DB Partition Number: 0
Sessions: 2
   2 Time: 20100930204152 Oldest log: S0000024.LOG DB Partition Number: 0
Sessions: 2
   3 Time: 20100930203906 Oldest log: S0000024.LOG DB Partition Number: 0
Sessions: 1
   4 Time: 20100930202923 Oldest log: S0000023.LOG DB Partition Number: 0
Sessions: 1
    5 Time: 20100930202350 Oldest log: S0000022.LOG DB Partition Number: 0
Sessions: 1
Retrieving INCREMENTAL DATABASE BACKUP information.
 No INCREMENTAL DATABASE BACKUP images found for TPCDB
Retrieving DELTA DATABASE BACKUP information.
  No DELTA DATABASE BACKUP images found for TPCDB
Retrieving TABLESPACE BACKUP information.
  No TABLESPACE BACKUP images found for TPCDB
Retrieving INCREMENTAL TABLESPACE BACKUP information.
  No INCREMENTAL TABLESPACE BACKUP images found for TPCDB
Retrieving DELTA TABLESPACE BACKUP information.
  No DELTA TABLESPACE BACKUP images found for TPCDB
Retrieving LOAD COPY information.
  No LOAD COPY images found for TPCDB
Retrieving LOG ARCHIVE information.
   Log file: S0000000.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-09-30-15.57.39
  Log file: S0000002.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-09-30-18.22.12
  Log file: S0000024.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-09-30-20.32.56
   Log file: S0000025.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-09-30-20.35.33
```

Figure 18-12 Output from a db2adutl query database TPCDB command

First, look at the "oldest log" number against the oldest backup version. After we deleted one backup as shown in Figure 18-11 on page 735, the oldest log is \$0000022.log.

Next, look at the list of log archive files from the same output to see if there are any earlier logs. If there are earlier logs and you do not want to wait for Tivoli Storage Manager to expire them, use the following command to delete them. See Figure 18-13.

db2adutl delete logs between S0000000 and S0000002 database TPCDB

Tip: When you specify log numbers, add the **S** at the start of the number but not the **.L0G** at the end.

```
[db2inst1@tpcblade6-11 ~]$ db2adutl delete logs between S0000000 and S0000002 database TPCDB

Query for database TPCDB

Retrieving LOG ARCHIVE information.

Log file: S0000000.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2010-09-30-15.57.39

Do you want to delete this log image (Y/N)? Y

Are you sure (Y/N)? Are you sure (Y/N)? Y

Log file: S0000002.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2010-09-30-18.22.12

Do you want to delete this log image (Y/N)? Y

Are you sure (Y/N)? Are you sure (Y/N)? Y
```

Figure 18-13 Example command to delete DB2 archive logs

18.11 Verifying a backup file

To ensure that a backup file is valid and you will be able to restore from it, you can use the **db2adut1** command to check the integrity of a specific backup, entering the backup datestamp and database name as parameters:

db2adutl verify full taken at 20100922172357 db TPCDB

A verification process starts, as shown in Figure 18-14, and the following confirmation message is issued:

Image Verification Complete - successful.

```
[db2inst1@tpcblade6-11 ~]$ db2adutl verify full taken at 20100930204152 db TPCDB
Query for database TPCDB
Retrieving FULL DATABASE BACKUP information. Please wait.
  FULL DATABASE BACKUP image:
      ./TPCDB.O.db2inst1.NODE0000.CATN0000.20100930204152.001, DB Partition Number: 0
      ./TPCDB.0.db2inst1.NODE0000.CATN0000.20100930204152.002, DB Partition Number: 0
  Do you wish to verify this image (Y/N)? Y
  Verifying file: ./TPCDB.O.db2inst1.NODE0000.CATN0000.20100930204152.001
Read O bytes, assuming we are at the end of the image
                     ./TPCDB.O.db2inst1.NODE0000.CATN0000.20100930204152.002
WARNING only partial image read, bytes read: 16384 of 16781312
Read O bytes, assuming we are at the end of the image
Image Verification Complete - successful.
Retrieving INCREMENTAL DATABASE BACKUP information.
 No INCREMENTAL DATABASE BACKUP images found for TPCDB
Retrieving DELTA DATABASE BACKUP information.
 No DELTA DATABASE BACKUP images found for TPCDB
```

Figure 18-14 Performing a backup verification

If the verification fails, that backup is not usable and you must take a new one.

18.12 Restoring Tivoli Storage Productivity Center database

This section describes the steps to restore the DB2 repository database for Tivoli Storage Productivity Center. As with the backup process, restoring from an online backup is more complex than restoring from an offline backup.

Restoring from an offline backup is a simple point-in-time exercise. Because the database was stopped at the time of the offline backup, it is logically consistent and you can restore the data "as is." However, circular logging does not offer the ability to roll forward through database changes using the logs to recover to an exact point-in-time. Therefore, if you take a database backup on a 24 hour cycle, you lose updates to the Tivoli Storage Productivity Center repository that were made between these points.

When you configure archive logging, you can restore a backup and then roll forward through the logs to any point-in-time to minimize data loss. This way gives you an enhanced level of protection to the Tivoli Storage Productivity Center repository data at the expense of more complexity in the process. You cannot simply restore a backup taken online as is, because an online backup is not logically consistent in its own right. Following an online restore, some roll forward is necessary to bring the restored database to a consistent and usable state.

Finally, we do not intend for this section to be a comprehensive guide to the DB2 restore commands. We intend to give you the basic restore functions that you need to recover a database from both file system and Tivoli Storage Manager backups. For details about this subject, see *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441.

18.12.1 Restoring from offline backups

Restoring from an offline backup is the simplest type of restore operation. It brings the database back to the specific point-in-time that the backup was taken. You can then restart Tivoli Storage Productivity Center.

Overview of basic steps

Restoring an offline backup from a file system or Tivoli Storage Manager involves the following basic steps:

- 1. Stop the Tivoli Storage Productivity Center services if they are still running.
- 2. Choose the backup image from which to restore.
- Restore the TPCDB database.
- 4. Restart the Tivoli Storage Productivity Center service.
- 5. Resolve potential agent issues after you restore. For more information, see 18.12.3, "Potential agent issues after the restore process" on page 747.

Stopping the Tivoli Storage Productivity Center services

Stop the Tivoli Storage Productivity Center services on Linux by using the commands in Example 18-19.

Example 18-19 Linux commands to stop Tivoli Storage Productivity Center

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 stop
/<usr or opt>/IBM/TPC/device/bin/linux/stopTPCF.sh

Choosing the backup image to restore from file system

If the backup image that you require is stored in Tivoli Storage Manager, skip this section and follow the steps in "Choosing a backup image to restore from Tivoli Storage Manager" on page 741.

Use one of the following steps:

- ► List the contents of the directory where you stored the backups, and choose a backup image from which to restore.
- ► Use the DB2 command list history backup all for TPCDB (in a DB2 command window) to see a list of the backup versions that are available.

Figure 18-15 shows the available backup files. For example, we can select the following file (which translates to a backup taken on 30 September 2010 at 14:13:35):

TPCDB.O.db2inst1.NODE0000.CATN0000.20100930141335.001

```
[db2inst1@tpcblade6-11 ~]$ ls -l /var/TPC_database_backups
total 2017944
-rw----- 1 db2inst1 db2iadm1 218177536 Sep 30 14:13
TPCDB.O.db2inst1.NODE0000.CATN0000.20100930141335.001
-rw----- 1 db2inst1 db2iadm1 234958848 Sep 30 17:34
TPCDB.O.db2inst1.NODE0000.CATN0000.20100930173433.001
-rw----- 1 db2inst1 db2iadm1 218177536 Sep 30 17:39
TPCDB.O.db2inst1.NODE0000.CATN0000.20100930173932.001
-rw----- 1 db2inst1 db2iadm1 218177536 Sep 30 17:41
TPCDB.O.db2inst1.NODE0000.CATN0000.20100930174149.001
-rw----- 1 db2inst1 db2iadm1 234958848 Sep 30 17:57
TPCDB.O.db2inst1.NODE0000.CATN0000.20100930175729.001
-rw----- 1 db2inst1 db2iadm1 234958848 Sep 30 19:00
TPCDB.O.db2inst1.NODE0000.CATN0000.20100930190048.001
-rw----- 1 db2inst1 db2iadm1 234958848 Sep 30 19:10
TPCDB.O.db2inst1.NODE0000.CATN0000.20100930191010.001
-rw----- 1 db2inst1 db2iadm1 234958848 Sep 30 19:10
TPCDB.O.db2inst1.NODE0000.CATN0000.20100930191031.001
-rw----- 1 db2inst1 db2iadm1 234958848 Sep 30 19:10
TPCDB.O.db2inst1.NODE0000.CATN0000.20100930191041.001
```

Figure 18-15 Viewing backup versions available for restore

From the file name, we extract the backup image timestamp, which in this case is as follows: 20100930141335

You need this timestamp number for the next step, "Restoring the TPCDB database (offline)" on page 742.

Choosing a backup image to restore from Tivoli Storage Manager

If you have chosen a backup image from the file system, skip this step and go to "Restoring the TPCDB database (offline)" on page 742.

To search for a backup image in Tivoli Storage Manager, use the **db2adut1** command: db2adut1 guery full database TPCDB

Figure 18-16 shows example output from the db2adut1 command for the TPCDB database.

```
[db2inst1@tpcblade6-11 ~]$ db2adutl query full database TPCDB
Query for database TPCDB
Retrieving FULL DATABASE BACKUP information.
   1 Time: 20100930204416 Oldest log: S0000025.LOG DB Partition Number: 0
Sessions: 2
   2 Time: 20100930204152 Oldest log: S0000024.LOG DB Partition Number: 0
Sessions: 2
   3 Time: 20100930203906 Oldest log: S0000024.LOG DB Partition Number: 0
Sessions: 1
   4 Time: 20100930202923 Oldest log: S0000023.LOG DB Partition Number: 0
Sessions: 1
   5 Time: 20100930202350 Oldest log: S0000022.LOG DB Partition Number: 0
Sessions: 1
Retrieving INCREMENTAL DATABASE BACKUP information.
 No INCREMENTAL DATABASE BACKUP images found for TPCDB
Retrieving DELTA DATABASE BACKUP information.
 No DELTA DATABASE BACKUP images found for TPCDB
```

Figure 18-16 Command db2adutl example to query backup versions available

From the list, select a backup timestamp, as in the following example:

20100930204416

You need this timestamp number for the next step.

Restoring the TPCDB database (offline)

Log in to your DB2 server and switch to the DB instance owner ID (usually db2inst1) or source the instance profile:

. /home/db2inst1/sqllib/db2profile

Then, initiate the DB2 command line processor:

db2

To restore from file system backups, issue the commands, which are shown in Example 18-20, in the DB2 command-line processor and use the timestamps that you have selected.

Example 18-20 Restore command from file system backups

restore database TPCDB from /var/TPC database backups taken at 20100930141335

If you restore from Tivoli Storage Manager, use the commands that are shown in Example 18-21.

Example 18-21 Restore command from Tivoli Storage Manager backups

restore database TPCDB use TSM taken at 20100922170304

Figure 18-17 shows an example of the restore process dialog for the TPCDB database restore process from a file system.

db2 => restore database TPCDB from /var/TPC_database_backups taken at 20100930141335

SQL2539W Warning! Restoring to an existing database that is the same as the backup image database. The database files will be deleted.

Do you want to continue ? (y/n) y

DB20000I The RESTORE DATABASE command completed successfully.

Figure 18-17 Example of offline restore of TPCDB from a file system

Restarting the Tivoli Storage Productivity Center services

After you restore the Tivoli Storage Productivity Center database, restart the Tivoli Storage Productivity Center services to bring the server back online. Issue the commands shown in Example 18-22.

Example 18-22 Windows commands to start Tivoli Storage Productivity Center

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 start
/<usr or opt>/IBM/TPC/device/bin/linux/startTPCF.sh

18.12.2 Restoring from online backups

Restoring from an online backup can be more complex than restoring from an offline backup, because there are more choices about what to do after you restore the backup image.

You might restore to a backup image from a week ago, because you actually want your Tivoli Storage Productivity Center environment put back to that point. You might want to restore from the last known good backup and roll forward through the archive logs to get your Tivoli Storage Productivity Center database as close as possible to the point before the problem occurred that triggered the need to restore.

Overview of basic steps

Restoring an online backup from a file system or Tivoli Storage Manager involves the following basic steps:

- 1. Stop the Tivoli Storage Productivity Center services if they are not already stopped.
- 2. Choose the backup image from which to restore.
- Restore the TPCDB database.
- 4. Roll forward the database.
- 5. Restart the Tivoli Storage Productivity Center services.
- 6. Resolve any new agent issues after you restore.

Stopping the Tivoli Storage Productivity Center services

Stop the Tivoli Storage Productivity Center services by using the commands in Example 18-23.

Example 18-23 Linux commands to stop Tivoli Storage Productivity Center

/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 stop
/<usr or opt>/IBM/TPC/device/bin/linux/stopTPCF.sh

Choosing the backup image from which to restore

Choose a backup image from which to restore by using the same process as you use for offline backups.

See "Choosing the backup image to restore from file system" on page 740 or "Choosing a backup image to restore from Tivoli Storage Manager" on page 741.

Restoring the TPCDB database (online)

The initial process of restoring a database that was taken online is the same as the offline process. However, when you complete the restoration, you are *not* ready to use the database. After you restore the backup, the database status is Roll-Forward Pending. "Rolling the database forward" on page 745 explains how to proceed from this point.

To restore the database, log in to your DB2 server and switch to the DB2 instance owner ID (usually db2inst1) or source the instance profile:

. /home/db2inst1/sqllib/db2profile

Then, initiate the DB2 command-line processor:

db2

To restore the database from file system backups, issue the commands in Example 18-24 in the DB2 command line processor using the timestamp that you have selected.

Example 18-24 Restore command from file system backups

restore database TPCDB from /var/TPC database backups taken at 20100930214725

If you restore from Tivoli Storage Manager, use commands as in Example 18-25.

Example 18-25 Restore command from Tivoli Storage Manager backups

restore database TPCDB use TSM taken at 20100930204416

Figure 18-18 shows an example of the restore process dialog for the TPCDB database restore from file system.

db2 => restore database TPCDB from /var/TPC_database_backups taken at 20100930214725

SQL2539W Warning! Restoring to an existing database that is the same as the

backup image database. The database files will be deleted. Do you want to continue ? (y/n) y

bo you want to continue; (y/n) y

DB20000I The RESTORE DATABASE command completed successfully.

Figure 18-18 Example of online restore of TPCDB from file system

Rolling the database forward

After the database restore processes complete, you can start the roll forward. You cannot start Tivoli Storage Productivity Center at this point, because the database will not open until you perform some type of roll forward.

Roll-forward options in DB2 can be complex. We do not intend to provide a complete guide to DB2 roll forward recovery. We describe how to roll forward in two ways:

► Roll forward to the end of the logs

This way rolls forward from the restore point through all available log files to the most recent consistent point-in-time. If you are using an old backup and there are many logs through which to roll, this method can take some time.

Roll forward to a point-in-time

With a point-in-time roll forward, you can specify a specific point-in-time for the roll forward process to stop, complete, and allow the database to open.

Rolling the database forward to the end of the logs

To roll the database forward to the end of all of the logs after a restore, type the following commands in the DB2 command-line processor (Figure 18-19):

rollforward database TPCDB to end of logs and complete

When each command completes, it returns an audit of the process.

Tip: The last committed transaction time is displayed in a UTC-0 time format even if your local time zone is, for example, PDT (UTC-8).

```
db2 => rollforward database TPCDB to end of logs and complete
                               Rollforward Status
Input database alias
                                      = TPCDB
Number of nodes have returned status
                                      = 1
Node number
                                      = 0
Rollforward status
                                      = not pending
Next log file to be read
Log files processed
                                     = S0000000.LOG - S0000000.LOG
                                    = 2010-10-01-04.47.28.000000 UTC
Last committed transaction
DB20000I The ROLLFORWARD command completed successfully.
```

Figure 18-19 Roll forward TPCDB to the end of the logs and complete

When complete, proceed to "Restarting the Tivoli Storage Productivity Center services" on page 746.

Rolling the database forward to a point-in-time

Tip: By default, DB2 uses UTC-0 time for the point-in-time roll forward. Add the **use local time** flag to the command if you want to specify a time in your local time zone.

Use the following steps to roll the database forward to a given point-in-time after the restore:

- Use the DB2 command line processor as seen in Figure 18-20 to enter the rollforward command. In this example, we rolled the TPCDB database forward to a few minutes after the restore time. We entered the time using the use local time option.
- 2. Enter the point-in-time in the YYYY-MM-DD-HH, MM, SS format.

The command for the TPCDB database is as follows, for example:

rollforward database TPCDB to 2010-09-30-22.40 using local time and complete

```
db2 => rollforward database TPCDB to 2010-09-30-22.40 using local time and
complete
                               Rollforward Status
Input database alias
                                     = TPCDB
Number of nodes have returned status
                                    = 1
Node number
                                     = 0
Rollforward status
                                     = not pending
Next log file to be read
Log files processed
                                    = S0000000.LOG - S0000000.LOG
Last committed transaction = 2010-09-30-22.35.52.000000 Local
DB20000I The ROLLFORWARD command completed successfully.
```

Figure 18-20 Roll forward the TPCDB to point-in-time and complete

Notice that the actual Last committed transaction time differs slightly from the time that is requested in the roll forward. This time is the closest that DB2 can get to the requested time and still keep the database in a consistent state.

Restarting the Tivoli Storage Productivity Center services

After you complete the restore operation and the roll forward for the Tivoli Storage Productivity Center database, restart the Tivoli Storage Productivity Center service to bring the server back online. To do this, issue the commands shown in Example 18-26.

Example 18-26 Linux commands to start Tivoli Storage Productivity Center

```
/<usr or opt>/IBM/TPC/data/server/tpcdsrv1 start
/<usr or opt>/IBM/TPC/device/bin/linux/startTPCF.sh
```

18.12.3 Potential agent issues after the restore process

Following the restore process of the Tivoli Storage Productivity Center database to a previous state, there is always a possibility that new agents were deployed to machines after the time of the restore process. If this happens, certain agents, which are not registered in the Tivoli Storage Productivity Center server and which are now running with an older version of the database, are out in the environment. To correct this problem, instruct the orphaned agents to re-register themselves with the Tivoli Storage Productivity Center server.

Reinstall the agents with the force parameter by using the Agent command or a deployment job from the GUI.

18.13 Backup scheduling and automation

This section provides information about the method to use to back up the database and how often to do it.

18.13.1 Frequency of full TPCDB backups

How often you take a full backup of your Tivoli Storage Productivity Center database depends on how critical the TPCDB data is to your business.

Run the full TPCDB backup once a week. If it is significantly critical, implement a TPCDB backup strategy to accommodate your business needs. For example, a full TPCDB backup can be scheduled every weekend, and incremental backups (not explained in this chapter) can be scheduled every week day.

For details about this subject, see IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference, SC27-2441.

18.13.2 TPCDB backup automation

Taking backups of the TPCDB database can be automated. Several available options for this task are as follows:

- ► Linux cron scheduler
- ▶ DB2 Administration Server's scheduler
- Tivoli Storage Manager Backup-Archive Scheduler.

See the following publications for more details:

- ▶ IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference, SC27-2441
- Backing Up DB2 Using IBM Tivoli Storage Management, SG24-6247



Tivoli Storage Productivity Center database backup on AIX

This chapter explains how to plan for backing up and restoring the Tivoli Storage Productivity Center database that resides in DB2 on the Tivoli Storage Productivity Center server in an AIX environment.

The chapter covers both offline backup (cold backup) and online backup (hot backup) of the database along with the merits of each type.

The Tivoli Storage Productivity Center product does not provide any extra backup and recovery tools in addition to those tools already provided with the DB2 product. This chapter is not intended to be a comprehensive guide to all functions of backup and recovery built into DB2. For details about this subject, see the *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441.

19.1 Before you start

We suggest that you read Chapter 17, "Tivoli Storage Productivity Center database considerations" on page 687 before you continue with this chapter. To improve the performance of your DB2 instance, tune your database settings by following the instructions that are provided in this section, and review 17.4.3, "Database backup method considerations" on page 711.

You can also estimate the storage requirements for the Tivoli Storage Productivity Center repository database, and plan the sizing of your backup volumes accordingly.

19.2 Common backup setup steps

Complete the following setup steps for both file system and Tivoli Storage Manager backups:

1. Configure the DB2 history file to keep the number of backup versions that you want to retain. Your organization might already have a policy for how many versions you need to keep.

Change the DB2 num_db_backups parameter and set the value to the number of backup versions that you require. You must also set the rec_his_retentn parameter to a value of -1. By setting this value to -1, the rec_his_retentn parameter follows the value set in the num_db_backups parameter.

Important: For this change to take effect, Tivoli Storage Productivity Center service must be stopped and started. This restarting does not necessarily need to happen directly after you change the parameter.

- 2. Log in to your DB2 server and switch to the DB instance owner ID (usually db2inst1) or source the instance profile:
 - . /home/db2inst1/sqllib/db2profile

Then, initiate the DB2 command-line processor:

db2

3. Example 19-1 shows how to set the num_db_backups value to 4 versions and rec_his_retentn to -1 for both the Tivoli Storage Productivity Center databases.

Issue the commands at the db2 => prompt in the command-line processor window.

Example 19-1 DB2 commands to configure how many backup versions to keep

```
connect to TPCDB
update db cfg using num_db_backups 4
update db cfg using rec_his_retentn -1
disconnect TPCDB
quit
```

Important: When you set new values for num_db_backups and rec_his_retentn, the new values are not effective until you stop all database connections.

 Restart Tivoli Storage Productivity Center to make the changes effective. You can either reboot the server, or alternatively stop and start the services, as shown in Example 19-2.

Example 19-2 AIX commands to stop and start Tivoli Storage Productivity Center services

```
stopsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/stopTPCF.sh

startsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/startTPCF.sh
```

19.3 Offline backup to file system setup steps

This section describes how to set up offline backup for the Tivoli Storage Productivity Center server database to flat files in a file system. Because the offline backup method is the default method for Tivoli Storage Productivity Center, there is little DB2 configuration needed before you can perform a backup.

Important: Ensure that you perform the steps in 19.2, "Common backup setup steps" on page 750 and these steps.

The steps are as follows:

Choose a location to use for the DB2 backup output. Choose a directory that has enough
free space to hold the number of backups that you plan to retain. It is best to use a
separate file system rather than the file system that contains the DB2 database.

You can choose to use a location that is a remotely mounted CIFS or NFS, so that the backup data is secured to another server, perhaps at another location in your organization.

This example uses the directory /var/TPC_database_backups.

Important: DB2 does not create this directory for you. Create this directory before you attempt a backup, and make sure that user db2inst1 has write permissions.

2. Create a batch script to control the backup process.

Two files are used:

- /home/root/TPCBKP/TPC_backup_offline_file

This file (shown in Example 19-3) runs the backup.

Example 19-3 File /home/root/TPCBKP/TPC_backup_offline_file

```
#!/bin/ksh

#This is a sample backup script
#To backup TPC offline
#To disk filesystems

. /home/db2inst1/sqllib/db2profile
echo "Stopping Tivoli Storage Productivity Center services"
echo "-----"
echo
stopsrc -s TSRMsrv1
```

```
/opt/IBM/TPC/device/bin/aix/stopTPCF.sh
echo
echo "Starting backup of the DB2 database"
echo "-----"
db2 force application all
db2 $(cat /home/root/TPCBKP/database_list_offline_file.txt )
echo
echo "Restarting Tivoli Storage Productivity Center services"
echo "-----"
echo
startsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/startTPCF.sh
echo
echo "Offline backup process complete"
echo "------"
exit 0
```

```
Remember: Make the script executable by using the chmod command:

chmod +x /home/root/TPCBKP/TPC_backup_offline_file
```

- /home/root/TPCBKP/database_list_offline_file.txt

This file (shown in Example 19-4) is DB2 scripted list of databases to back up.

Example 19-4 File /home/root/TPCBKP/database_list_offline_file.txt

backup database TPCDB to /var/TPC database backups without prompting

See 19.7.1, "Performing an offline backup to a file system" on page 762 to run an offline backup.

19.4 Offline backup to Tivoli Storage Manager setup steps

This section describes the steps necessary to set up an offline backup of the Tivoli Storage Productivity Center server database to a Tivoli Storage Manager server. The backup to Tivoli Storage Manager is a little more complex to set up but does not require you to set aside large amounts of local disk space for backup versions on the Tivoli Storage Productivity Center server.

In this section, we assume the following requirements:

- You have a basic working knowledge of Tivoli Storage Manager.
- ► An operational Tivoli Storage Manager server already exists to which you can send backup data.
- ➤ Your Tivoli Storage Manager administrator has defined storage, which will receive the backups, to the policies.
- ➤ You have already installed a Tivoli Storage Manager Backup-Archive client on the Tivoli Storage Productivity Center server, and you have configured it to do standard file backups.
- ➤ You have installed the Tivoli Storage Manager API Client on the Tivoli Storage Productivity Center server.
- You used default installation paths for Tivoli Storage Manager.

Stop and restart: You must stop Tivoli Storage Productivity Center and DB2 as part of this configuration process. Plan this exercise at a time when you can restart Tivoli Storage Productivity Center.

Use the following steps, described in this section, to configure DB2 to Tivoli Storage Manager integration:

- 1. "Adding new variables to AIX" on page 753
- 2. "Configuring Tivoli Storage Manager option file and password" on page 754
- 3. "Restarting DB2" on page 755
- 4. "Creating an offline backup to Tivoli Storage Manager script" on page 755

19.4.1 Adding new variables to AIX

Table 19-1 shows a list of Tivoli Storage Manager API environment variables to add to AIX. The listed values assume a default installation of Tivoli Storage Manager on the Tivoli Storage Productivity Center server.

Table 19-1 System environment variables

Environment variable name	Value
DSMI_DIR	/usr/tivoli/tsm/client/api/bin64
DSMI_CONFIG	/home/db2inst1/tsm/dsm.opt
DSMI_LOG	/home/db2inst1/tsm

Set the environment variables for the API client. As Example 19-5 shows, add the following lines in the \$HOME/.profile file for the DB2 instance administrator (usually /home/db2inst1/.profile):

Example 19-5 Add Tivoli Storage Manager variables to the DB2 profile

```
echo 'export DSMI_DIR=/usr/tivoli/tsm/client/api/bin64
export DSMI_CONFIG=/home/db2inst1/tsm/dsm.opt
export DSMI_LOG=/home/db2inst1/tsm
' >> /home/db2inst1/sqllib/db2profile
```

Important:

- ► If you are using a 32-bit version of DB2, use /opt/tivoli/tsm/client/api/bin location instead of /opt/tivoli/tsm/client/api/bin64 location.
- If it does not exist, create a /home/db2inst1/tsm directory and make sure that user db2inst1 is the owner.

19.4.2 Configuring Tivoli Storage Manager option file and password

This section describes the steps necessary to configure the Tivoli Storage Manager option file dsm.opt and then set the Tivoli Storage Manager password so that the DB2 backup process can communicate with the Tivoli Storage Manager API.

Important: At this stage, your Tivoli Storage Manager client must be already registered with a Tivoli Storage Manager server. If the Tivoli Storage Manager server is accepting open registrations, just by starting the Tivoli Storage Manager client GUI or command line you will be asked for a password to register your client. If the Tivoli Storage Manager server is using closed registration, you will need the Tivoli Storage Manager administrator to register your client.

The steps are as follows:

1. Edit the dsm.sys file, which is in the following location:

/usr/tivoli/tsm/client/api/bin64/

32-bit version: If you are using a 32-bit version of DB2, edit the following file instead: /usr/tivoli/tsm/client/api/bin/dsm.sys

Make sure that the client option PASSWORDACCESS is set to GENERATE as shown in Figure 19-1.

```
tpcblade4-14v3> cat /usr/tivoli/tsm/client/api/bin64/dsm.sys

Servername TSMsrv1
COMMMethod TCPIP
TCPPort 1500
TCPSERVERADDRESS tsmsrv1.storage.tucson.ibm.com
PASSWORDACCESS GENERATE
```

Figure 19-1 Contents of the dsm.sys file

2. Create or edit the dsm.opt file, which is in the following location:

```
/home/db2inst1/tsm/
```

The dsm.opt file requires having only one line in it, which is reference to the server stanza in the dsm.sys file. In our case, it is TSMsrv, as shown in Figure 19-2.

```
tpcblade4-14v3> cat /home/db2inst1/tsm/dsm.opt
Servername TSMsrv1
```

Figure 19-2 Contents of the dsm.opt file

- 3. Set the Tivoli Storage Manager password so that DB2 can authenticate with the Tivoli Storage Manager server when DB2 performs a backup or restore operation:
 - a. Run the dsmapipw command as shown in Figure 19-3.
 - b. Enter the current and new Tivoli Storage Manager password. You can reuse the existing Tivoli Storage Manager password.

Important: You must run the dsmapipw command even if you do not intend to change the Tivoli Storage Manager password. Running this command registers the password with the Tivoli Storage Manager API. Registering this password in the setup phase means that a DB2 operator can perform backup and restore operations without needing to know the Tivoli Storage Manager client password. If a Tivoli Storage Manager administrator changes or resets the Tivoli Storage Manager password, you need to run the dsmapipw command again.

Figure 19-3 Running the dsmapipw command

Important: Check that files dsierror.log and dsm.opt in the /home/db2inst1/tsm directory are owned by the DB2 instance owner (db2inst1) to avoid errors during the backup process.

19.4.3 Restarting DB2

Now that you have completed the configuration steps, restart Tivoli Storage Productivity Center and DB2 as shown in Example 19-6 to ensure that the environment variables are picked up.

Example 19-6 AIX commands to restart Tivoli Storage Productivity Center services and DB2

```
. /home/db2inst1/sqllib/db2profile

stopsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/stopTPCF.sh

db2 force applications all
db2stop
db2start

startsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/startTPCF.sh
```

19.4.4 Creating an offline backup to Tivoli Storage Manager script

Create two files:

► The first file is the script that you run (Example 19-7): /home/root/TPCBKP/TPC backup offline tsm

```
#!/bin/ksh
#This is a sample backup script
#To backup TPC offline
#To Tivoli Storage Manager
. /home/db2inst1/sqllib/db2profile
echo "Stopping Tivoli Storage Productivity Center services"
echo "-----"
echo
stopsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/stopTPCF.sh
echo
echo "Starting backup of the DB2 database"
echo "-----"
db2 force application all
db2 $(cat /home/root/TPCBKP/database list offline tsm.txt )
echo "Restarting Tivoli Storage Productivity Center services"
startsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/startTPCF.sh
echo "Offline backup process complete"
echo "-----"
exit 0
```

```
Remember: Make the script executable by using the chmod command:

chmod +x /root/TPCBKP/TPC_backup_offline_tsm
```

The second file is the DB2 scripted list of databases to back up (Example 19-8):

```
/home/root/TPCBKP/database_list_offline_tsm.txt
```

Example 19-8 File /home/root/TPCBKP/database_list_offline_tsm.txt

backup database TPCDB use tsm without prompting

19.5 Online backup to Tivoli Storage Manager setup steps

This section describes the steps for configuring the Tivoli Storage Productivity Center database to enable for online backup to Tivoli Storage Manager. The significant difference between online and offline backup is the need to enable archive logging on the database. As we describe in 17.4.3, "Database backup method considerations" on page 711, operating in this mode provides many backup and recovery benefits at the expense of increased complexity in the database operation.

Important: You must stop Tivoli Storage Productivity Center to perform these tasks:

- ▶ DB2 requires a full backup of each database before you can start the Tivoli Storage Productivity Center database again after these reconfiguration steps. We include the instructions to perform a full backup of the database. Allow time in your outage planning for the backup to complete.
- ► Also, complete the steps in 19.2, "Common backup setup steps" on page 750 to set the number of backup versions that you want to retain in the history file.

Be sure that you consider the advantages and disadvantages of archive logging before you continue with this setup. For full details of DB2 logging methods, see the DB2 product manuals. Also see *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441, for detailed information about this subject.

Considerations:

- ▶ If you set up DB2 for online backup to Tivoli Storage Manager, you cannot easily change to an online backup to file system. You must choose between these methods, because you are setting the destination for the archive logging process. If you decide in the future to change to the online file system method, you will need to reconfigure DB2 to send the archive logs to file system. To complete the task, this reconfiguration requires restarting Tivoli Storage Productivity Center.
- ▶ It is possible to perform an online backup to file system and have the archive logs going to Tivoli Storage Manager. However, this approach is not desirable because the difficulty of managing and tracking information makes this method a poor practice.

Set up and test DB2 to Tivoli Storage Manager integration before you attempt the steps in this section. Use 19.4, "Offline backup to Tivoli Storage Manager setup steps" on page 752. When you are satisfied that DB2 is communicating with Tivoli Storage Manager and you have performed at least one successful offline backup, return to this section.

19.5.1 DB2 parameter changes for archive logging to Tivoli Storage Manager

To set up archive logging to Tivoli Storage Manager, complete the following tasks:

1. You must make a number of parameter choices for the configuration of archive logging as shown in Table 19-2. These parameters determine where DB2 keeps its log files. Be sure that the db2inst1 user is the owner of all log directories.

Table 19-2 DB2 parameters

DB2 parameter	Example value	Comment
Primary log path	/var/DB2/active_logs	This location is where DB2 keeps the current logs for the database. For best performance, place these logs on a separate volume than the volume that holds the data.
Failed log path	/var/DB2/failed_logs	This location is where DB2 put log files if the archive process fails. This can happen if Tivoli Storage Manager is down or unreachable when DB2 tries to send a log file to Tivoli Storage Manager.

2. Stop Tivoli Storage Productivity Center by using the commands in Example 19-9.

Example 19-9 AIX commands to stop Tivoli Storage Productivity Center

stopsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/stopTPCF.sh

- 3. Log in to your DB2 server and switch to the DB2 instance owner ID (usually db2inst1) or source the instance profile:
 - . /home/db2inst1/sqllib/db2profile

Then, initiate the DB2 command-line processor: db2

4. Issue the commands from Example 19-10 in the command-line processor window. Substitute your chosen values for the parameters that form part of the UPDATE DB CFG command. See Table 19-2 on page 757. The final command performs an offline backup of the database.

Important: The database backup is required after this reconfiguration, and the DB2 database will not open again until the database backup is completed.

Example 19-10 DB2 command to configure archive logging to Tivoli Storage Manager

CONNECT TO TPCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR TPCDB USING logarchmeth1 TSM failarchpath
/var/DB2/failed_logs newlogpath /var/DB2/active_logs

BACKUP DATABASE TPCDB USE TSM

QUIT

Attention: Verify that the following directories exist and are owned by the db2inst1 user:

/var/DB2/logs/DB2_failed_logs
/var/DB2/logs/DB2 active logs

5. When the database backups is complete, you can restart Tivoli Storage Productivity Center. Issue the commands shown in Example 19-11.

Example 19-11 Start Tivoli Storage Productivity Center

startsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/startTPCF.sh

19.5.2 Creating an online backup script for Tivoli Storage Manager

Create two files:

► The script (Example 19-12) that you run to start the backup: /home/root/TPCBKP/TPC backup online tsm

Example 19-12 File /home/root/TPCBKP/TPC_backup_online_tsm

```
#!/bin/ksh

#This is a sample backup script
#To backup TPC online
#To Tivoli Storage Manager

. /home/db2inst1/sqllib/db2profile
echo
echo "Starting backup of the DB2 database"
echo "-----"
db2 $(cat /home/root/TPCBKP/database_list_online_tsm.txt)
echo
echo "Online backup process complete"
echo "-----"
exit 0
```

```
Remember: Make the script executable by using the chmod command:

chmod +x /home/root/TPCBKP/TPC backup online tsm
```

► The DB2 scripted list of databases (Example 19-13) to back up:

```
/home/root/TPCBKP/database_list_online_tsm.txt
```

Example 19-13 File C:\scripts\database_list_online_tsm.txt

backup database TPCDB online use tsm without prompting

19.6 Online backup to a file system setup steps

Performing online backups to a file system requires you to set up archive logging to a file system also. When operating with this method, DB2 does not clean up old and no longer necessary archive log files. Therefore, you need to put processes in place to clean up old log files after a specific amount of time to prevent the system from filling up. You also need to plan for this amount of space. Over a number of weeks, the log space that is required for a Tivoli Storage Productivity Center database can become many times larger than the database.

To be able to restore an online DB2 database taken two weeks ago, for example, you must have log files going back to that same date that you can restore. An online DB2 database backup is not stand-alone, because you cannot restore the online DB2 database backup without at least some logs for it to roll forward to a consistent state.

Important:

- Although switching between a backup destination of online to file system and online to Tivoli Storage Manager is a straightforward process, switching the logging path is not as easy. To switch the logging from Tivoli Storage Manager to a file system requires a stop and a start of the database and, therefore, a restart of the Tivoli Storage Productivity Center services.
- Choose either a Tivoli Storage Manager backup or a file system backup, and stay with that specific method.

19.6.1 Setting up DB2 archive logging to a file system

Set up DB2 archive logging to a file system by using the following steps:

 To configure archive logging, you must make several parameter choices. The parameters in Table 19-3 determine where DB2 will keep its log files, how many log files to keep, and the size of the log files. Be sure that the db2inst1 user is the owner of all log directories.

Table 10-3	DB2 parameters t	or archive	logging to	a fila evetam
iable 19-3	DDZ Darameters i	or archive i	ioaairia lo	a ille system

DB2 parameter	Example value	Comment
Primary log path	/var/DB2/active_logs	The location where DB2 will keep the current logs for the database. For best performance, place the logs on a separate volume than the data.
Archive log path	/var/DB2/archive_logs/TPCDB	The location where DB2 will archive log files for the TPCDB database.
Failed log path	/var/DB2/failed_logs	The location where DB2 will put log files if the archive process fails, which can happen if the file system for the primary logs fills up. Choose a location that is <i>not</i> on the same file system as the archive logs.

2. Choose a file system path (Table 19-4) to store the DB2 database backups. Ensure that that the directory is owned by user db2inst1.

Table 19-4 File system location for database backups

Database backup path
/var/TPC_database_backups

3. Stop Tivoli Storage Productivity Center by using the commands in Example 19-14.

Example 19-14 AIX commands to stop Tivoli Storage Productivity Center

stopsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/stopTPCF.sh

- 4. Log in to your DB2 server and switch to the DB instance owner ID (usually "db2inst1") or source the instance profile:
 - . /home/db2inst1/sqllib/db2profile

Next, initiate the DB2 command-line processor:

db2

 Issue the commands shown in Example 19-15 on page 761 in the command-line processor window. Substitute your values for the parameters that form part of the UPDATE DB CFG command. See Table 19-3. Note that the final command performs an offline backup of the database.

Important: The offline backup of the database is required after the reconfiguration; the DB2 database does not open until the backup is complete.

Example 19-15 DB2 command to configure archive logging to a file system

```
CONNECT TO TPCDB

QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS

UNQUIESCE DATABASE

CONNECT RESET

UPDATE DB CFG FOR TPCDB USING logarchmeth1 DISK:/var/DB2/archive_logs/TPCDB failarchpath/var/DB2/failed_logs newlogpath /var/DB2/active_logs

BACKUP DATABASE TPCDB TO /var/TPC_database_backups
```

6. When both database backups complete, you can restart Tivoli Storage Productivity Center. Issue the commands shown in Example 19-16.

```
Example 19-16 Start Tivoli Storage Productivity Center
```

```
startsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/startTPCF.sh
```

19.6.2 Creating online backup script to file system

Create two files to control the backup process for online backup to file system output:

► The script (Example 19-17) that you run to start the backup:

```
/home/root/TPCBKP/TPC_backup_online_file
```

Example 19-17 File /home/root/TPCBKP/TPC_backup_online_file

```
#!/bin/ksh

#This is a sample backup script

#To backup TPC online

#To Tivoli Storage Manager

. /home/db2inst1/sqllib/db2profile
echo
echo "Starting backup of the DB2 database"
echo "-----"
db2 $(cat /home/root/TPCBKP/database_list_online_file.txt)
echo
echo "Online backup process complete"
echo "------"
exit 0
```

Remember: Make the script executable by using the chmod command:

chmod +x /home/root/TPCBKP/TPC backup online file

► The DB2 scripted list of databases to back up (Example 19-18):

/home/root/TPCBKP/database list online file.txt

Example 19-18 File /home/root/database_list_online_file.txt

backup database TPCDB online to /var/TPC_database_backups without prompting

19.7 Performing offline database backups

This section describes how to perform the offline backup of the Tivoli Storage Productivity Center database.

Running an offline DB2 database backup takes Tivoli Storage Productivity Center out of service for the period of the backup. This impacts data collections from CIMOMs, and you might miss other infrastructure events.

Make sure that you understand the impact of stopping Tivoli Storage Productivity Center in your environment before proceeding. If your environment cannot tolerate regular stoppages of Tivoli Storage Productivity Center for a backup operation, consider configuring and using online backup.

19.7.1 Performing an offline backup to a file system

Important: You must complete the initial steps as detailed in 19.2, "Common backup setup steps" on page 750 before you can start offline backups.

To do an offline backup to a file system, run the **TPC_backup_offline_file** script (Figure 19-4). This script stops the Tivoli Storage Productivity Center process, performs a backup of the DB2 database to the location specified in the scripts, and then restarts the services.

```
p55ap1(root)/> /home/root/TPCBKP/TPC backup offline file
Stopping Tivoli Storage Productivity Center services
-----
0513-044 The TSRMsrv1 Subsystem was requested to stop.
Setting Variables for SANM
Stopping server1 with default options
ADMU0116I: Tool information is being logged in file
   /opt/IBM/TPC/device/apps/was/profiles/deviceServer/logs/server1/stopS
erver.log
ADMU0128I: Starting tool with the deviceServer profile
ADMU3100I: Reading configuration for server: server1
ADMU3201I: Server stop request issued. Waiting for stop status.
ADMU4000I: Server server1 stop completed.
Starting backup of the DB2 database
DB20000I The FORCE APPLICATION command completed successfully.
DB21024I This command is asynchronous and may not be effective immediately.
Backup successful. The timestamp for this backup image is: 20101001144033
Restarting Tivoli Storage Productivity Center services
-----
0513-059 The TSRMsrv1 Subsystem has been started. Subsystem PID is 1130594.
Setting Variables for SANM
Starting server1 for Device Manager
ADMU0116I: Tool information is being logged in file
   /opt/IBM/TPC/device/apps/was/profiles/deviceServer/logs/server1/start
Server.log
ADMU0128I: Starting tool with the deviceServer profile
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 508002
Offline backup process complete
-----
```

Figure 19-4 Running an offline backup to a file system

19.7.2 Performing an offline backup to Tivoli Storage Manager

Important: You must complete the initial setup steps that are described in 19.4, "Offline backup to Tivoli Storage Manager setup steps" on page 752 before you can start offline backups.

Running an offline DB2 database backup takes Tivoli Storage Productivity Center out of service for the period of the backup. Make sure it is acceptable to take Tivoli Storage Productivity Center out of service before you proceed.

To perform an offline backup to Tivoli Storage Manager, run the TPC_backup_offline_tsm script as shown in Figure 19-5.

```
p55ap1(root)/> /home/root/TPCBKP/TPC backup offline tsm
Stopping Tivoli Storage Productivity Center services
0513-044 The TSRMsrv1 Subsystem was requested to stop.
Setting Variables for SANM
Stopping server1 with default options
ADMU0116I: Tool information is being logged in file
/opt/IBM/TPC/device/apps/was/profiles/deviceServer/logs/server1/stopServer.log
ADMU0128I: Starting tool with the deviceServer profile
ADMU3100I: Reading configuration for server: server1
ADMU3201I: Server stop request issued. Waiting for stop status.
ADMU4000I: Server server1 stop completed.
Starting backup of the DB2 database
DB20000I The FORCE APPLICATION command completed successfully.
DB21024I This command is asynchronous and may not be effective immediately.
Backup successful. The timestamp for this backup image is: 20101001144758
Restarting Tivoli Storage Productivity Center services
_____
0513-059 The TSRMsrv1 Subsystem has been started. Subsystem PID is 1028106.
Setting Variables for SANM
Starting server1 for Device Manager
ADMU0116I: Tool information is being logged in file
/opt/IBM/TPC/device/apps/was/profiles/deviceServer/logs/server1/startServer.log
ADMU0128I: Starting tool with the deviceServer profile
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 331802
Offline backup process complete
-----
```

Figure 19-5 Running an offline backup to Tivoli Storage Manager

19.8 Performing online database backup

This section describes how to run online backups of the Tivoli Storage Productivity Center database. By running the backups online (or hot), it is not necessary to stop the Tivoli Storage Productivity Center service, which means that the backup process does not interrupt processes, such as performance data collection.

Important:

- Before you can perform online DB2 backup, you must configure Tivoli Storage Manager to DB2 communication and DB2 archive logging to Tivoli Storage Manager or DB2 to file system logging.
- ▶ If you plan to use online backup to Tivoli Storage Manager, perform the setup steps in 19.4, "Offline backup to Tivoli Storage Manager setup steps" on page 752 and 19.5, "Online backup to Tivoli Storage Manager setup steps" on page 756.

19.8.1 Performing an online database backup to Tivoli Storage Manager

This section describes how to run an online backup of the Tivoli Storage Productivity Center database to Tivoli Storage Manager. You must have already run the appropriate setup steps.

To do a backup, run the TPC_backup_online_tsm script as shown in Figure 19-6.

```
p55ap1(root)/> /home/root/TPCBKP/TPC_backup_online_tsm

Starting backup of the DB2 database
------

Backup successful. The timestamp for this backup image is : 20101001144937

Online backup process complete
```

Figure 19-6 Running an online backup to Tivoli Storage Manager

19.8.2 Performing an online backup to a file system

This section describes how to run an online backup of the Tivoli Storage Productivity Center database to file system output files. You must already have completed the setup steps described in 19.6, "Online backup to a file system setup steps" on page 759.

To do the backup, run the TPC backup online file script as shown in Figure 19-7.

```
p55ap1(root)/> /home/root/TPCBKP/TPC_backup_online_file

Starting backup of the DB2 database
------

Backup successful. The timestamp for this backup image is : 20101001145102

Online backup process complete
-----------
```

Figure 19-7 Running an online backup to file system

19.9 Other backup considerations

Apart from the DB2 database, a number of important files and directories must be backed up up to preserve the state of a Tivoli Storage Productivity Center server installation. Back up all files under the Tivoli Storage Productivity Center installation directory. In addition, the Tivoli GUID is stored in the /etc/Tivoli/TSRM location. Be sure to also preserve this Tivoli GUID.

The important Tivoli Storage Productivity Center server directories to secure are as follows:

- ▶ TPC_Server_install_dir/config/
- TPC_Server_install_dir/data/config/
- ► TPC Server install dir/device/conf/

These directories contain the various configuration files for your installation. Saving these directories is important, because they might be customized configurations and not the default configurations.

19.10 Managing database backup versions

This section is an overview of the maintenance processes for which you need to plan. DB2 does not prune older versions automatically. Therefore, you must maintain the number of Tivoli Storage Productivity Center database backup versions that you want on a file system or Tivoli Storage Manager.

19.10.1 Managing backup versions for a file system

This section describes what you need to know to manage DB2 backups that were performed to disk.

Deleting unwanted backups

DB2 does not automatically manage the deletion of the unwanted database backups or archive logs from the file system. You need to create a maintenance plan to delete the old backups. If you plan to use online backup to a file system, you need to create a plan to delete the old archive logs.

How DB2 organizes backups on a file system

When you do a backup to a file system, you supply the backup script with a path to use. Since the release of Version 9.1, DB2 creates backup files with a file name containing a datestamp that details when the backup was taken.

Figure 19-8 shows the backup directory of the TPCDB database. The directory contains backups taken at separate times on 01 October 2010. DB2 timestamps all backups in this way; every time a backup is made, a new file is created:

- The first part of the file name starts as follows:
 - TPCDB.O.DB2.NODE0000.CATN0000
- ► The last part of the file name consists of the date in yyyyMMDD format, as in the following example:

20100930

Plan to delete old backup files to suit the requirements of your backup and recovery policy.

```
p55ap1(root)/> ls /var/TPC_database_backups
TPCDB.0.db2inst1.NODE0000.CATN0000.20101001143814.001
TPCDB.0.db2inst1.NODE0000.CATN0000.20101001145102.001
TPCDB.0.db2inst1.NODE0000.CATN0000.20101001144033.001
```

Figure 19-8 DB2 backup directory

19.10.2 Managing archive log files on a file system

If you plan to do online backups, you must configure DB2 to use archive logging. If you plan to do online backups to disk, you must also maintain the archive logs directory on a regular basis.

Figure 19-9 shows the directory structure for logging on to the TPCDB database. Over time, this directory fills up with logs. If your recovery policy is to keep backup versions for five days, you must keep logs in this directory for at least the same period of time, because you cannot restore an online backup without logs from the same date and time to make the recovery valid.

Notice that the directory that holds the logs is named C0000000, which is the log cycle number. If you restore the database, the cycle number increments by one and starts in C0000001 and so on. Ensure that any automated deletion process that you implement can handle this numbering.

```
p55ap1(root)/> ls -R /var/DB2/archive_logs/TPCDB/

/var/DB2/archive_logs/TPCDB/:
db2inst1

/var/DB2/archive_logs/TPCDB/db2inst1:
TPCDB

/var/DB2/archive_logs/TPCDB/db2inst1/TPCDB:
NODE0000

/var/DB2/archive_logs/TPCDB/db2inst1/TPCDB/NODE0000:
C0000000

/var/DB2/archive_logs/TPCDB/db2inst1/TPCDB/NODE0000/C00000000:
S0000001.LOG S0000002.LOG S0000003.LOG S0000004.LOG S0000005.LOG S0000006.LOG
```

Figure 19-9 DB2 archive logs

19.10.3 Managing backup versions that you store in Tivoli Storage Manager

This section describes how to maintain, view, and delete backup data and archive logs that you have sent to Tivoli Storage Manager. DB2 does not automatically prune backup versions and log files from Tivoli Storage Manager. You need to use the db2adut1 tool to perform these housekeeping functions.

Tip: This section is not intended to be a comprehensive guide to the **db2adut1** tool. The intent here is to detail the commands that you likely will use on a regular basis to maintain the data that is held in Tivoli Storage Manager.

Purpose of the db2adutl command

The command line tool, db2adut1, communicates with Tivoli Storage Manager through its API interface. Use this tool to interrogate the backup and archive log data that is stored in Tivoli Storage Manager at any one time, verify that you no longer require old backups, and delete unnecessary old backups.

Reasons to use db2adutl

When DB2 stores a backup session in Tivoli Storage Manager, DB2 always stores the backup session with a unique file name, which is the timestamp when the backup was made. This means that these backup versions never get superseded by a new version with the same file name. The backup files remain "active" versions in Tivoli Storage Manager, and, therefore, Tivoli Storage Manager never deletes the backup versions. Use the command, <code>db2adut1</code>, to select unwanted backup versions and tell Tivoli Storage Manager to flag them as "inactive." In this way, Tivoli Storage Manager then deletes them over time based on the standard policy rules that the Tivoli Storage Manager administrator sets.

You handle DB2 archive logs differently. They are stored in Tivoli Storage Manager as "archive" data, which means Tivoli Storage Manager retains them for a set period of time based on its policies. Although you can use db2adut1 to explicitly remove DB2 archive logs, if Tivoli Storage Manager archive retention policy is set appropriately, this way is not necessary.

Important: Make sure that the Tivoli Storage Manager archive retention policy that you use to store the DB2 logs is set for a sufficient period of time to allow recovery of your oldest database backup. However, you also want to make sure that the policy for the retention period is not so long that it wastes storage space in Tivoli Storage Manager.

How to query backups held in Tivoli Storage Manager

Next, we explain how to guery backups that are held in Tivoli Storage Manager.

Important: The **db2adut1** command is normally in the bin directory of the DB2 instance owner, which is usually the following location:

/home/db2inst1/sqllib/bin/db2adutl location.

Before using the command, log in with the db2inst1 user ID.

The db2adut1 command to query database backup versions are as follows:

► The following command lists all the database versions and the logs that are held for all databases stored in Tivoli Storage Manager and TPCDB, in this case:

```
db2adut1 query
```

► The following command lists all database versions and logs for the TPCDB database. Note that the database name is case sensitive and is in capital letters:

```
db2adutl query database TPCDB
```

Figure 19-10 shows the sample output from this command. The output shows that two database backups are stored in Tivoli Storage Manager and six archive logs.

► The following command has a shorter output. It lists only the database backup versions and the archive logs:

```
db2adutl query full
```

```
$ db2adut1 query database TPCDB
Query for database TPCDB
Retrieving FULL DATABASE BACKUP information.
    1 Time: 20101001150356 Oldest log: S0000007.LOG DB Partition Number: 0
Sessions: 2
    2 Time: 20101001150044 Oldest log: S0000006.LOG DB Partition Number: 0
Sessions: 2
    3 Time: 20101001145751 Oldest log: S0000005.LOG DB Partition Number: 0
Sessions: 2
    4 Time: 20101001145519 Oldest log: S0000004.LOG DB Partition Number: 0
Sessions: 2
    5 Time: 20101001144937 Oldest log: S0000002.LOG DB Partition Number: 0
Sessions: 2
    6 Time: 20101001144758 Oldest log: S0000002.LOG DB Partition Number: 0
Sessions: 1
    7 Time: 20101001142657 Oldest log: S0000000.LOG DB Partition Number: 0
Sessions: 1
-- 8< ---- OUTPUT CLIPPED -- 8< ----
Retrieving LOG ARCHIVE information.
   Log file: S0000000.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.40.15
   Log file: S0000001.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.47.38
   Log file: S0000002.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.49.53
   Log file: S0000003.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.51.02
   Log file: S0000004.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.55.20
   Log file: S0000005.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.57.52
   Log file: S0000006.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-15.00.46
   Log file: S0000007.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-15.03.58
```

Figure 19-10 Sample output from a db2adutl query database TPCDB command

Deleting backup versions held in Tivoli Storage Manager

The following commands and examples show how to delete database backup versions that are held in Tivoli Storage Manager:

Important: Because Tivoli Storage Manager does not allow the root user to delete backups that are created by the db2inst1 instance, log in with user ID db2inst1 before trying these commands.

► The following command deletes backup versions from Tivoli Storage Manager that are older than three days. This type of command is useful, because you can easily script it to run each day to remove the older backup.

```
db2adutl delete full older than 90 days
You can also specify a database name:
db2adutl delete full older than 90 days database TPCDB
```

Figure 19-11 is an example of running this command.

► The following command deletes all backup versions from Tivoli Storage Manager, except for the last three versions. Again, this command is useful when scripting an automatic

process.

db2adutl delete full keep 5
You can also specify a database name:
db2adutl delete full keep 5 database TPCDB

```
$ db2adutl delete full keep 5 database TPCDB
Query for database TPCDB
Retrieving FULL DATABASE BACKUP information.
 Taken at: 20101001144758 DB Partition Number: 0
                                                      Sessions: 1
 Taken at: 20101001142657 DB Partition Number: 0
                                                      Sessions: 1
    Do you want to delete these backup images (Y/N)? Y
     Are you sure (Y/N)? Y
The current delete transaction failed. You do not have
sufficient authorization. Attempting to deactivate
backup image(s) instead...
Success.
Retrieving INCREMENTAL DATABASE BACKUP information.
 No INCREMENTAL DATABASE BACKUP images found for TPCDB
Retrieving DELTA DATABASE BACKUP information.
 No DELTA DATABASE BACKUP images found for TPCDB
```

Figure 19-11 Example of a db2adutl delete full keep 5 database TPCDB command

Deleting DB2 archive log files in Tivoli Storage Manager

The commands described in this section are examples of how to delete database archive logs from Tivoli Storage Manager.

Important:

- ▶ Be careful when you delete archive log files. If you delete logs that are still needed for some of your backup versions, you render those backups useless.
- Archive logs only exist in Tivoli Storage Manager if you have configured archive logging so that online backup is possible.
- ► Ask the Tivoli Storage Manager administrator to configure Tivoli Storage Manager to delete the archive logs on a regular basis by configuring the Tivoli Storage Manager "archive copy group" that DB2 uses. Set a retention period that suits your needs. If you use a general purpose archive copy group, Tivoli Storage Manager might keep all archive logs for several years causing unnecessary usage of the storage in your Tivoli Storage Manager environment.

To delete archive logs, first query the Tivoli Storage Manager server to establish which logs you want to delete.

To query the Tivoli Storage Manager server for the TPCDB database, issue the following command (see Figure 19-12):

db2adutl query database TPCDB

```
$ db2adut1 query database TPCDB
Query for database TPCDB
Retrieving FULL DATABASE BACKUP information.
    1 Time: 20101001150356 Oldest log: S0000007.LOG DB Partition Number: 0
    2 Time: 20101001150044 Oldest log: S0000006.LOG DB Partition Number: 0
Sessions: 2
    3 Time: 20101001145751 Oldest log: S0000005.LOG DB Partition Number: 0
Sessions: 2
    4 Time: 20101001145519 Oldest log: S0000004.LOG DB Partition Number: 0
Sessions: 2
    5 Time: 20101001144937 Oldest log: S0000002.LOG DB Partition Number: 0
Sessions: 2
-- 8< ---- OUTPUT CLIPPED -- 8< ----
Retrieving LOG ARCHIVE information.
   Log file: S0000000.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.40.15
   Log file: S0000001.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.47.38
   Log file: S0000002.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.49.53
   Log file: S0000003.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.51.02
   Log file: S0000004.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.55.20
   Log file: S0000005.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-14.57.52
   Log file: S0000006.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-15.00.46
   Log file: S0000007.LOG, Chain Num: 0, DB Partition Number: 0, Taken at:
2010-10-01-15.03.58
```

Figure 19-12 Output from a db2adutl query database TPCDB command

First, look at the "oldest log" number against the oldest backup version. After we delete one backup, as shown in Figure 19-11 on page 770, the oldest log is S0000002.log.

Next, look at the list of log archive files from the same output to determine whether any earlier logs exist. If earlier logs exist and you do not want to wait for Tivoli Storage Manager to expire them, use the following command to delete them. See Figure 19-13.

db2adutl delete logs between S0000000 and S0000001 database TPCDB

Tip: When you specify log numbers, add the **S** at the start of the number, but not the **.LOG** at the end.

```
$ db2adutl delete logs between $0000000 and $0000001 database TPCDB

Query for database TPCDB

Retrieving LOG ARCHIVE information.

Log file: $0000000.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2010-10-01-14.40.15

Do you want to delete this log image (Y/N)? Y

Are you sure (Y/N)? Y

Log file: $0000001.LOG, Chain Num: 0, DB Partition Number: 0, Taken at: 2010-10-01-14.47.38

Do you want to delete this log image (Y/N)? Y

Are you sure (Y/N)? Y
```

Figure 19-13 Example command to delete DB2 archive logs

19.11 Verifying a backup file

To ensure that a backup file is valid and you are able to restore from it, you can use the **db2adut1** command to check the integrity of a specific backup, entering the backup datestamp and database name as parameters:

```
db2adutl verify full taken at 20101001145751 db TPCDB
```

A verification process starts, as shown in Figure 19-14, and the following confirmation message is issued:

Image Verification Complete - successful.

```
$ db2adutl verify full taken at 20101001145519 db TPCDB
Query for database TPCDB
Retrieving FULL DATABASE BACKUP information. Please wait.
  FULL DATABASE BACKUP image:
     ./TPCDB.O.db2inst1.NODE0000.CATN0000.20101001145519.001, DB Partition Number: 0
     ./TPCDB.O.db2inst1.NODE0000.CATN0000.20101001145519.002, DB Partition Number: 0
  Do you wish to verify this image (Y/N)? Y
  Verifying file: ./TPCDB.O.db2inst1.NODE0000.CATN0000.20101001145519.001
Read O bytes, assuming we are at the end of the image
                   ./TPCDB.O.db2inst1.NODE0000.CATN0000.20101001145519.002
WARNING only partial image read, bytes read: 16384 of 1576960
Read O bytes, assuming we are at the end of the image
Image Verification Complete - successful.
Retrieving INCREMENTAL DATABASE BACKUP information.
 No INCREMENTAL DATABASE BACKUP images found for TPCDB
Retrieving DELTA DATABASE BACKUP information.
 No DELTA DATABASE BACKUP images found for TPCDB
```

Figure 19-14 Performing a backup verification

If the verification fails, that backup is not usable and you must take a new backup.

19.12 Restoring Tivoli Storage Productivity Center database

This section describes the steps to restore the DB2 repository database for Tivoli Storage Productivity Center. As with the backup process, restoring from an online backup is more complex than restoring from an offline backup.

Restoring from an offline backup is a simple point-in-time exercise. Because the database was stopped at the time of the offline backup, it is logically consistent and you can restore the data "as is." However, circular logging does not offer the ability to roll forward through database changes using the logs to recover to an exact point-in-time. Therefore, if you take a database backup on a 24 hour cycle, you lose updates to the Tivoli Storage Productivity Center repository that were made between these points.

When you configure archive logging, you can restore a backup and then roll forward through the logs to any point-in-time to minimize data loss. This way gives you an enhanced level of protection to the Tivoli Storage Productivity Center repository data at the expense of more complexity in the process. You cannot simply restore a backup taken online as is, because an online backup is not logically consistent in its own right. Following an online restore, some roll forward is necessary to bring the restored database to a consistent and usable state.

Finally, we do not intend for this section to be a comprehensive guide to the DB2 restore commands. We intend to give you the basic restore functions that you need to recover a database from both file system and Tivoli Storage Manager backups. For details about this subject, see *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441.

19.12.1 Restoring from offline backups

Restoring from an offline backup is the simplest type of restore operation. It brings the database back to the specific point-in-time that the backup was taken. You can then restart Tivoli Storage Productivity Center.

Overview of basic steps

Restoring an offline backup from a file system or Tivoli Storage Manager involves the following basic steps:

- 1. Stop the Tivoli Storage Productivity Center services if they are still running.
- 2. Choose the backup image from which to restore.
- 3. Restore the TPCDB database.
- 4. Restart the Tivoli Storage Productivity Center service.
- 5. Resolve potential agent issues after you restore. For more information, see 19.12.3, "Potential agent issues after the restore process" on page 781.

Stopping the Tivoli Storage Productivity Center services

Stop the Tivoli Storage Productivity Center services on AIX by using the commands in Example 19-19.

Example 19-19 AIX commands to stop Tivoli Storage Productivity Center

stopsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/stopTPCF.sh

Choosing the backup image to restore from file system

If the backup image that you require is stored in Tivoli Storage Manager, skip this section and follow the steps in "Choosing a backup image to restore from Tivoli Storage Manager" on page 776.

Use one of the following steps:

- ► List the contents of the directory where you stored the backups and choose a backup image from which to restore,
- ► Use the DB2 command list history backup all for TPCDB (in a DB2 command window) to see a list of the backup versions that are available.

Figure 19-15 shows the available backup files. For example, we can select the following file (which translates to a backup taken on 01 October 2010 at 14:38:14):

TPCDB.O.db2inst1.NODE0000.CATN0000.20101001143814.001

```
$ 1s -1 /var/TPC_database_backups
total 868096
-rw----- 1 db2inst1 db2iadm1 149676032 Oct 01 14:38
TPCDB.0.db2inst1.NODE0000.CATN0000.20101001143814.001
-rw----- 1 db2inst1 db2iadm1 149676032 Oct 01 14:40
TPCDB.0.db2inst1.NODE0000.CATN0000.20101001144033.001
-rw----- 1 db2inst1 db2iadm1 145100800 Oct 01 14:51
TPCDB.0.db2inst1.NODE0000.CATN0000.20101001145102.001
```

Figure 19-15 Viewing backup versions available for restore

From the file name, we extract the backup image timestamp, which in this case is as follows: 20101001143814

You need this timestamp number for the next step, "Restoring the TPCDB database (offline)" on page 777.

Choosing a backup image to restore from Tivoli Storage Manager

If you have chosen a backup image from the file system, ignore this step and go to "Restoring the TPCDB database (offline)" on page 777.

To search for a backup image in Tivoli Storage Manager, use the db2adut1 command: db2adut1 query full database TPCDB

Figure 19-16 shows example output from the db2adut1 command for the TPCDB database.

```
$ db2adut1 query full database TPCDB
Query for database TPCDB
Retrieving FULL DATABASE BACKUP information.
   1 Time: 20101001150356 Oldest log: S0000007.LOG DB Partition Number: 0
Sessions: 2
   2 Time: 20101001150044 Oldest log: S0000006.LOG DB Partition Number: 0
Sessions: 2
   3 Time: 20101001145751 Oldest log: S0000005.LOG DB Partition Number: 0
Sessions: 2
    4 Time: 20101001145519 Oldest log: S0000004.LOG DB Partition Number: 0
Sessions: 2
   5 Time: 20101001144937 Oldest log: S0000002.LOG DB Partition Number: 0
Sessions: 2
Retrieving INCREMENTAL DATABASE BACKUP information.
 No INCREMENTAL DATABASE BACKUP images found for TPCDB
Retrieving DELTA DATABASE BACKUP information.
 No DELTA DATABASE BACKUP images found for TPCDB
```

Figure 19-16 Command db2adutl example to query backup versions available

From the list, select a backup timestamp, as in the following example:

20101001145519

You need this timestamp number for the next step.

Restoring the TPCDB database (offline)

Log in to your DB2 server and switch to the DB instance owner ID (usually db2inst1) or source the instance profile:

. /home/db2inst1/sqllib/db2profile

Then, initiate the DB2 command-line processor:

db2

To restore from file system backups, issue the commands, which are shown in Example 19-20, in the DB2 command-line processor and use the timestamps that you selected.

Example 19-20 Restore command from file system backups

restore database TPCDB from /var/TPC_database_backups taken at 20101001143814

If you restore from Tivoli Storage Manager, use the commands that are shown in Example 19-21.

Example 19-21 Restore command from Tivoli Storage Manager backups

restore database TPCDB use TSM taken at 20101001145519

Figure 19-17 shows an example of the restore process dialog for the TPCDB database restore process from Tivoli Storage Manager.

db2 => restore database TPCDB use TSM taken at 20101001145519 SQL2539W Warning! Restoring to an existing database that is the same as the backup image database. The database files will be deleted. Do you want to continue ? (y/n) Y DB20000I The RESTORE DATABASE command completed successfully.

Figure 19-17 Offline restoring of TPCDB from Tivoli Storage Manager

Restarting the Tivoli Storage Productivity Center services

After you restore the Tivoli Storage Productivity Center database, restart the Tivoli Storage Productivity Center services to bring the server back online. Issue the commands shown in Example 19-22.

Example 19-22 AIX commands to start Tivoli Storage Productivity Center

startsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/startTPCF.sh

19.12.2 Restoring from online backups

Restoring from an online backup can be more complex than restoring from an offline backup, because there are more choices about what to do after you restore the backup image.

You might restore to a backup image from a week ago, because you actually want your Tivoli Storage Productivity Center environment put back to that point. You might want to restore from the last known good backup and roll forward through the archive logs to get your Tivoli Storage Productivity Center database as close as possible to the point before the problem occurred that triggered the need to restore.

Overview of basic steps

Restoring an online backup from a file system or Tivoli Storage Manager involves the following basic steps:

- Stop the Tivoli Storage Productivity Center services if they are not already stopped.
- 2. Choose the backup image from which to restore.
- 3. Restore the TPCDB database.
- Roll forward the database.
- 5. Restart the Tivoli Storage Productivity Center services.
- 6. Resolve any new agent issues after you restore.

Stopping the Tivoli Storage Productivity Center services

Stop the Tivoli Storage Productivity Center services by using the commands in Example 19-23.

Example 19-23 AIX commands to stop Tivoli Storage Productivity Center

stopsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/stopTPCF.sh

Choosing the backup image from which to restore

Choose a backup image from which to restore by using the same process as you use for offline backups.

See "Choosing the backup image to restore from file system" on page 775 or "Choosing a backup image to restore from Tivoli Storage Manager" on page 776.

Restoring the TPCDB database (online)

The initial process of restoring a database that was taken online is the same as the offline process. However, when you complete the restoration, you are *not* ready to use the database. After you restore the backup, the database status is Roll-Forward Pending. "Rolling the database forward" on page 779 explains how to proceed from this point.

To restore the database, log in to your DB2 server and switch to the DB instance owner ID (usually db2inst1) or source the instance profile:

. /home/db2inst1/sqllib/db2profile

Then, initiate the DB2 command line processor:

db2

To restore the database from file system backups, issue the commands in Example 19-24 in the DB2 command line processor using the timestamp that you have selected.

Example 19-24 Restore command from file system backups

restore database TPCDB from /var/TPC database backups taken at 20101001145102

If you restore from Tivoli Storage Manager, use commands as in Example 19-25.

Example 19-25 Restore command from Tivoli Storage Manager backups

restore database TPCDB use TSM taken at 20101001150044

Figure 19-18 shows an example of the restore process dialog for the TPCDB database restore from file system.

db2 => restore database TPCDB from /var/TPC_database_backups taken at 20101001145102
SQL2539W Warning! Restoring to an existing database that is the same as the backup image database. The database files will be deleted.
Do you want to continue ? (y/n) y
DB20000I The RESTORE DATABASE command completed successfully.

Figure 19-18 Example of online restore of TPCDB from file system

Rolling the database forward

After the database restore processes complete, you can start the roll forward. You cannot start Tivoli Storage Productivity Center at this point, because the database will not open until you perform some type of roll forward.

Roll-forward options in DB2 can be complex. We do not intend to provide a complete guide to DB2 roll forward recovery. We describe how to roll forward in two ways:

Roll forward to the end of the logs:

This way rolls forward from the restore point through all available log files to the most recent consistent point-in-time. If you are using an old backup and there are many logs through which to roll, this method can take some time.

► Roll forward to a point-in-time:

With a point-in-time roll forward, you can specify a specific point-in-time for the roll forward process to stop, complete, and allow the database to open.

Rolling the database forward to the end of the logs

To roll the database forward to the end of all of the logs after a restoration, type the following commands in the DB2 command-line processor (Figure 19-19):

rollforward database TPCDB to end of logs and complete

When each command completes, it returns an audit of the process.

Tip: The last committed transaction time is displayed in a UTC-0 time format even if your local time zone is, for example, PDT (UTC-8).

Figure 19-19 Roll forward TPCDB to the end of the logs and complete

When complete, proceed to "Restarting the Tivoli Storage Productivity Center services" on page 781.

Rolling the database forward to a point-in-time

Tip: By default, DB2 uses UTC-0 time for the point-in-time roll forward. Add the **use local time** flag to the command if you want to specify a time in your local time zone.

Use the following steps to roll the database forward to a given point-in-time after the restore:

- Use the DB2 command-line processor as shown in Figure 19-20 to enter the rollforward command. In this example, we rolled the TPCDB database forward to a few minutes after the restore time. We entered the time using the use local time option.
- 2. Enter the point-in-time in the YYYY-MM-DD-HH.MM.SS format.

The command for the TPCDB database is as follows, for example: rollforward database TPCDB to 2010-09-30-22.40 using local time and complete

```
db2 => rollforward database TPCDB to 2010-10-01-15.00 using local time and
complete
                                Rollforward Status
Input database alias
                                       = TPCDB
Number of nodes have returned status
                                      = 1
Node number
                                       = 0
Rollforward status
                                       = not pending
Next log file to be read
Log files processed
                                      = S0000004.LOG - S0000006.LOG
Last committed transaction
                                     = 2010-10-01-14.59.01.000000 Local
DB20000I The ROLLFORWARD command completed successfully.
```

Figure 19-20 Roll forward the TPCDB to point-in-time and complete

Notice that the actual last committed transaction time differs slightly from the time that is requested in the roll forward. This time is the closest that DB2 can get to the requested time and still keep the database in a consistent state.

Restarting the Tivoli Storage Productivity Center services

After you complete the restore operation and the roll forward for the Tivoli Storage Productivity Center database, restart the Tivoli Storage Productivity Center service to bring the server back online. To do this, issue the commands shown in Example 19-26.

Example 19-26 AIX commands to start Tivoli Storage Productivity Center

startsrc -s TSRMsrv1
/opt/IBM/TPC/device/bin/aix/startTPCF.sh

19.12.3 Potential agent issues after the restore process

Following the restore process of the Tivoli Storage Productivity Center database to a previous state, there is always a possibility that new agents were deployed to machines after the time of the restore process. If this happens, certain agents, which are not registered in the Tivoli Storage Productivity Center server and which are now running with an older version of the database, are out in the environment. To correct this problem, instruct the orphaned agents to re-register themselves with the Tivoli Storage Productivity Center server.

Reinstall the agents with the **force** parameter by using the **Agent** command or a deployment job from the GUI.

19.13 Backup scheduling and automation

This section provides information about the method to use to back up the database and how often to do it.

19.13.1 Frequency of full TPCDB backups

How often you take a full backup of your Tivoli Storage Productivity Center database depends on how critical the TPCDB data is to your business.

Run the full TPCDB backup once a week. If it is significantly critical, implement a TPCDB backup strategy to accommodate your business needs. For example, a full TPCDB backup can be scheduled every weekend, and incremental backups (not explained in this chapter) can be scheduled every week day.

For more details, see *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference*, SC27-2441.

19.13.2 TPCDB backup automation

Taking backups of the TPCDB database can be automated. Several available options for this task are as follows (see Figure 19-21):

- ► AIX cron scheduler
- ▶ DB2 Administration Server's scheduler
- ► Tivoli Storage Manager Backup-Archive Scheduler

See the following publications for more details:

- ► IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference, SC27-2441
- ▶ Backing Up DB2 Using IBM Tivoli Storage Management, SG24-6247

```
[root@tpcblade6-11 ~]# ls -R /var/DB2/archive_logs/TPCDB/
/var/DB2/archive_logs/TPCDB/:
db2inst1

/var/DB2/archive_logs/TPCDB/db2inst1:
TPCDB

/var/DB2/archive_logs/TPCDB/db2inst1/TPCDB:
NODE0000

/var/DB2/archive_logs/TPCDB/db2inst1/TPCDB/NODE0000:
C0000000

/var/DB2/archive_logs/TPCDB/db2inst1/TPCDB/NODE0000/C0000000:
S0000001.LOG S00000002.LOG S0000003.LOG S0000004.LOG S0000005.LOG S0000006.LOG
```

Figure 19-21 DB2 archive logs

Lessons learned

In this chapter, we offer information that we encountered while writing this book that might help you. It includes support references and helpful tips.

20.1 User-defined properties for Fabrics and Switches

In contrast to the Disk Manager \rightarrow Storage Subsystems panel where you can click the magnifying glass next to a subsystem to open a panel with details about the selected device, this way is not possible from fabrics in the Fabric Manager \rightarrow Fabrics panel.

For subsystems, this way is helpful because you can also enter a label and user-defined properties for the device. To enter the same type of information for Fabrics and Switches, go to the Level view in the Topology Viewer, right-click the device, and select **Launch Detail Panel** (Figure 20-1).

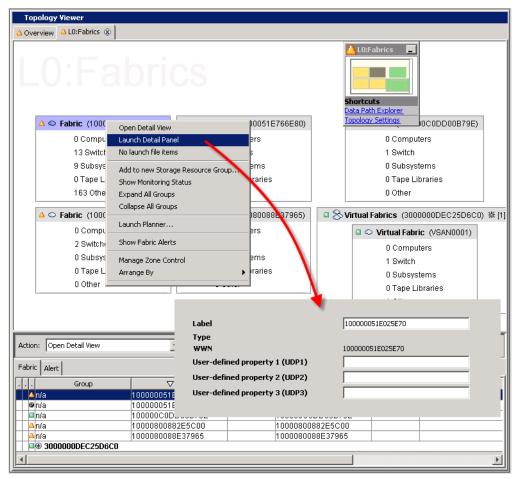


Figure 20-1 Setting Fabric user-defined properties

Providing your own labels, either as entirely new values or appending to existing values, can be especially helpful in certain cases. One example is Cisco fabrics where you might have VSANs named the same across various switches and want to be able to distinguish them within the Tivoli Storage Productivity Center GUI.

20.2 IBM Software Support Lifecycle

To determine whether your Tivoli Storage Productivity Center Version is still supported or whether you should upgrade, see the IBM Software Lifecyle website.

http://www.ibm.com/software/support/lifecycle/index t.html

Either browse through the long list of all products starting with the letter T, or use your browser's search function to look for the product name or product ID (PID). The column on the right shows when the product was released and when the support will end.

Select the check box for your products, and then click **View details** at the bottom of the page. A list of only the selected products opens, to which you can subscribe using RSS technology.

20.3 IBM Support Assistant

IBM Support Assistant V4 product add-ons are available for Tivoli Storage Productivity Center products. To download the IBM Support Assistant workbench code and unlock the details of how IBM Support Assistant can help you support Tivoli Storage Productivity Center, go to the following website:

 $\label{lem:http://www.ibm.com/support/docview.wss?rs=40&context=SSBSEX&context=SSMN28&context=SSMMUP\&context=SSBJFM\&dc=D600\&uid=swg21305569\&loc=en_US\&cs=utf-8\&lang=en$

20.4 Certificate errors in Windows Internet Explorer

Often when you work with Tivoli Storage Productivity Center or the web GUI of IBM devices, you will receive notifications about a certificate conflict because you are trying to establish a secure connection.

Most often, the two reasons for the notifications are shown at the top of the message that is being displayed (Figure 20-2):

- ► The certificate that is passed along from the website that you want to access is not yet stored in your local certificate store so it is a new certificate.
- ► The certificate that is passed along from the website that you want to access has a different IP address stored than the website that you are trying to access.

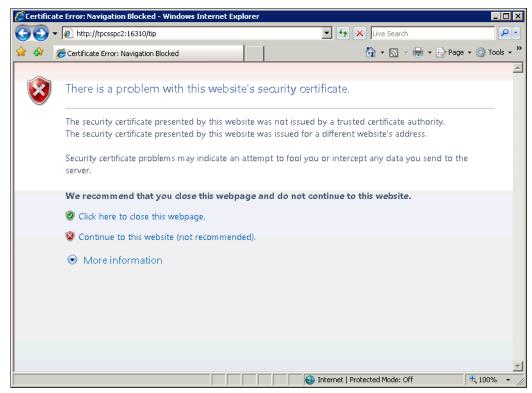


Figure 20-2 Certificate error message

The two messages usually occur only once. The second problem (message) must be fixed in the Internet Explorer configuration only once, but the first problem (message) is issued for any new website that provides a certificate that is not yet stored in your certificate store.

When you see the messages, complete the following steps:

- 1. "Preventing address certificate mismatch errors" on page 787
- "Managing new certificates" on page 788

Tip: You can resolve the certificate errors in Mozilla Firefox using similar steps.

20.4.1 Preventing address certificate mismatch errors

Configure Internet Explorer so that address mismatch messages are not issued as follows:

- 1. Open the Internet Explorer Options and go to the Advanced panel.
- Scroll to the security options and clear the Warn about certificate address mismatch* check box (Figure 20-3).

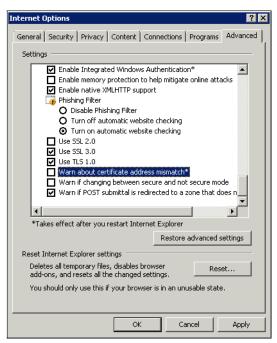


Figure 20-3 Warn about certificate address mismatch

3. Restart Internet Explorer for the changes to take effect.

20.4.2 Managing new certificates

Manage your new certificates and store them on your computer as follows:

- 1. Launch a web page, for example, the element manager for DS8000 or the login page for Tivoli Integrated Portal where you have seen this problem before (see Figure 20-2 on page 786).
- 2. Click the Continue to this website (not recommended) link.
- 3. At the top of the page, click **Certificate Error** (Figure 20-4).

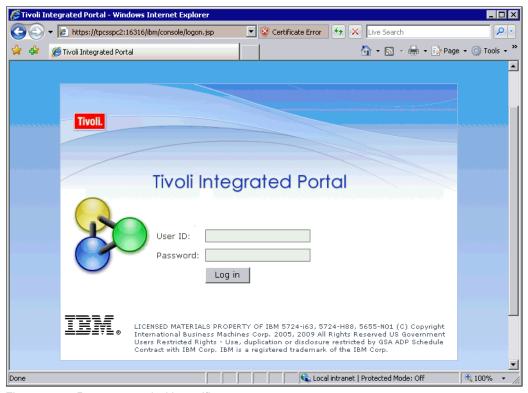


Figure 20-4 Pages opened with certificate error

4. The Certificate Invalid panel opens (Figure 20-5). Click View certificates.



Figure 20-5 View certificates

5. The Certificate dialog box opens (Figure 20-6). Click **Install Certificate**.



Figure 20-6 Certificate information

6. The Certificate Import Wizard opens (Figure 20-7). Click Next.



Figure 20-7 Certificate Import Wizard: Welcome panel

7. In the next panel (Figure 20-8), select **Place all certificates in the following store**, and then click **Browse**.



Figure 20-8 Certificate Import Wizard: Certificate Store

In the Select Certificate Store window (Figure 20-9), go to Trusted Root Certification
 Authorities → Local Computer, select the Show physical stores check box, and click OK.



Figure 20-9 Select Certificate Store panel

- 9. You are returned to the Import Certificate Wizard (Figure 20-8).Click Next.
- 10. In the Completing the Import Certificate Wizard window (Figure 20-10), click **Finish**.



Figure 20-10 Completing the Certificate Import Wizard

- 11. At the import was successful message, click **OK**.
- 12. Click **OK** to close the Certificate window.
- 13. Close the browser and re-launch the website that you tried to access.

20.5 Tivoli Storage Productivity Center support matrix

To find the Tivoli Storage Productivity Center support matrix for various components and versions, use the following steps:

- Go to the Tivoli Storage Productivity Center support website: http://www-947.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage Productivity Center Standard Edition
- Enter the words find support matrix in the field at the top right of the page. See Figure 20-11 for an example (when you open the link, the content of the page might differ from this example).

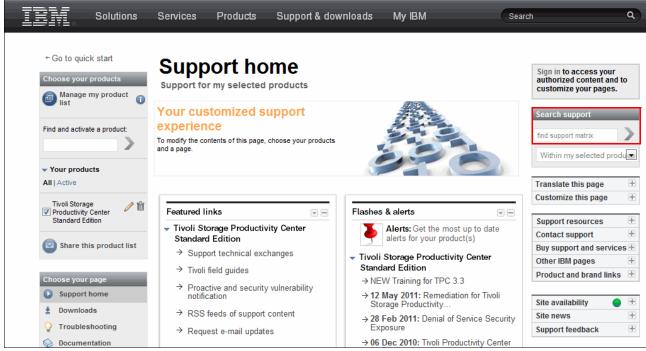


Figure 20-11 Search for "find support matrix"

3. Locate the document named "Find the Supported Products and Platforms Interoperability Matrix Links" (Figure 20-12), and open it.

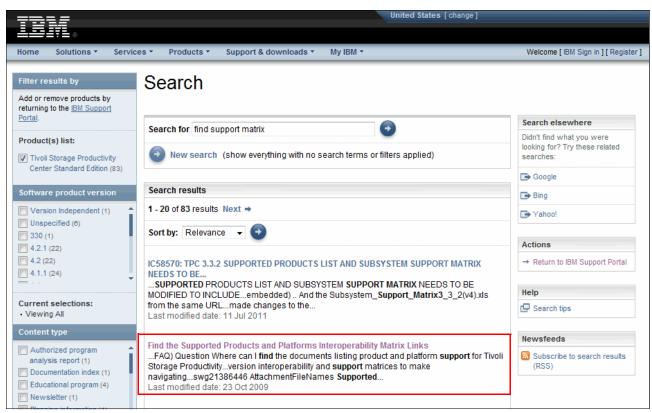


Figure 20-12 "Find the Supported Products and Platforms Interoperability Matrix Links"

The website contains tables that include links to the individual support matrices separated by version. See Figure 20-13.

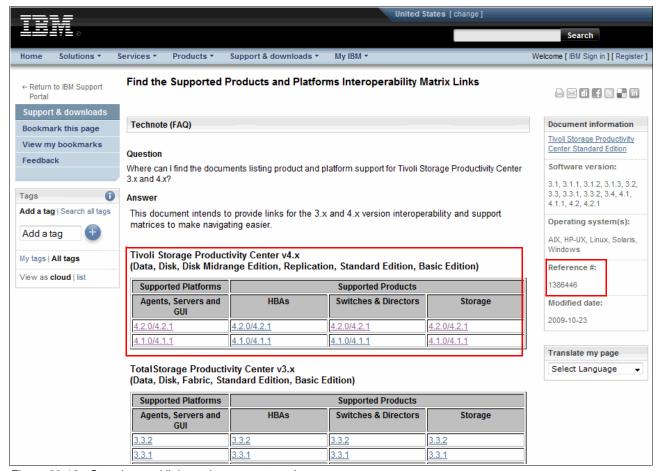


Figure 20-13 Overview and links to the support matrix

Tip: The document reference number, 1386446, is highlighted in Figure 20-13. Each document on the support site has a unique document number that becomes part of its website address. You can also search for a specific document using its reference number in the search field.

Throughout the support matrix documents, there are references to particular fix packs (for example, 4.2.1 FP1). As you work through this information, you might wonder how these fix packs correlate to the version displayed in the About panel on your server. You might also be curious whether there are newer fix packs available and how to download them.

Another support document is available to assist. It is "Latest Downloads for Tivoli Storage Productivity Center" (reference number 1320822). See Figure 20-14 for an example of the document. It has links at the top for finding the latest download for each of the releases. It also contains tables listing all of the field releases that allow you to correlate the fix packs and versions.

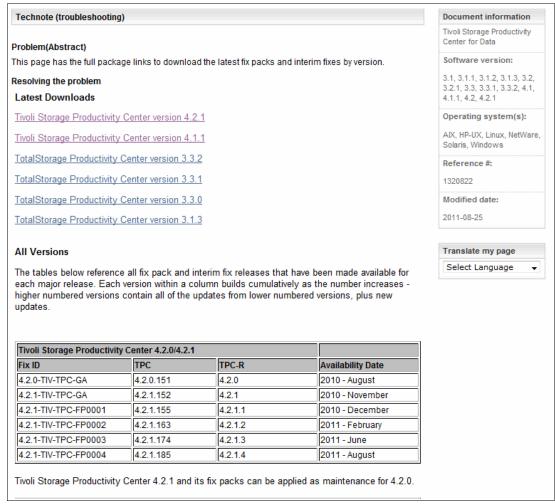


Figure 20-14 Latest Downloads for Tivoli Storage Productivity Center

20.6 DB2 hints

This section lists the two common DB2 errors:

- ► SQL5005C System Error
- User ID to stop and start DB2

20.6.1 SQL5005C System Error

This error indicates that the system tried to start the DB2 tray application, but the currently logged on user is not part of the DB2ADMNS group.

20.6.2 User ID to stop and start DB2

Only the user specified during the installation is able or allowed to start and stop DB2.

20.7 Extracting the installation images on AIX

Use the GNU tar program Version 1.14 or later rather than the native AIX tar program to extract files from Tivoli Storage Productivity Center installation images. The AIX tar program might truncate long file names, which can cause errors in Tivoli Storage Productivity Center and Tivoli Integrated Portal installation.

The GNU tar file archiving program can be found from the following link.

http://www-03.ibm.com/systems/power/software/aix/linux/toolbox/alpha.html

Additionally, you can ensure that the files extracted take on the User and Group ID values of the user running the program rather than those in the archive by using the -o flag.

Example 20-1 Extract the installation images using the -o flag

tar -xvof TPC_4.2.1.174_NL_aix_disk1_part1.tar

20.8 Tivoli Common Reporting hints

This section lists several precautions to keep in mind.

20.8.1 Linux libraries

When installing Tivoli Common Reporting (or Tivoli Storage Productivity Center for Replication) on Linux, it is vital to have the required libraries and packages present BEFORE attempting to install. You can use the following command on Linux to check the list of installed libraries.

```
rpm -qa | sort -d >/tmp/rpm.log
ls /lib >/tmp/lib.log
ls /usr/lib >/tmp/usr-lib.log
```

Examine the resulting log files from these commands to check for required libraries as documented in the *Tivoli Common Reporting Installation and Upgrade Guide*.

20.8.2 Physical Address Extension on Windows

In a 32-bit Windows environments, you can run into challenges with limited memory available. Even if your machine physically has more than 4 GB of RAM installed, you might be limited by having only 4 GB available for programs to use.

To overcome this limitation, make sure to enable Physical Address Extension (PAE). The PAE flag allows 32-bit machines to use more than 4 GB of RAM. To enable PAE on Windows 2008, use the bcdedit.exe tool. Running bcdedit.exe without options will show you the current setup. Run bcdedit /set pae ForceEnable to enable PAE. A reboot is required for these settings to take effect.

Tip: One way to make sure that you are not impacted by memory constraints is to use the bginfo.exe tool that is available from Microsoft. It will show detailed system information directly on your system desktop wallpaper. If you face a memory constrained setup, bginfo.exe will display 4 GB (or less) of memory available, even when the system properties shows that 16 GB of physical memory are installed.

20.8.3 Browser back button

When navigating in Tivoli Common Reporting, do not use the browser's back button to return to a previous screen. It might not provide the desired result. Rather, there is a back arrow in the upper right hand corner of the Tivoli Common Reporting interface that should be used.

20.9 Hints and tips

This chapter addresses only a subset of the hints and tips information that we have accumulated regarding Tivoli Storage Productivity Center and the products with which it interacts. We often refer to information documented in the Tivoli Storage Productivity Center Hints and Tips. This document is periodically updated and can be found by searching the Tivoli Storage Productivity Center support site with the words hints and tips. See 20.5, "Tivoli Storage Productivity Center support matrix" on page 792, for guidance on navigating the support site, if needed.





DB2 table space considerations

This appendix describes the trade-offs to consider when you determine which type of table space to use to store your data:

- System managed space (SMS)
- Database managed space (DMS)

Selecting an SMS or DMS table space

There are a number of trade-offs to consider when you determine which type of table space can best store your data. A table space can be managed using either system managed space (SMS), or database managed space (DMS):

- ► For an SMS table space, each container is a directory in the file space of the operating system, and the operating system's file manager controls the storage space.
- ► For a DMS table space, each container is either a fixed size pre-allocated file, or a physical device such as a disk, and the database manager controls the storage space.

Tables that contain user data exist in regular table spaces. The system catalog tables exist in a regular table space.

Tables that contain long field data or large object data, such as multimedia objects, exist in large table spaces or in regular table spaces. The base column data for these columns is stored in a regular table space; the long field or large object data can be stored in the same regular table space or in a specified large table space. Indexes can be stored in regular table spaces or large table spaces.

Temporary table spaces are classified as either system or user:

- System-temporary table spaces are used to store internal temporary data required during SQL operations such as sorting, reorganizing tables, creating indexes, and joining tables. Although you can create any number of system temporary table spaces, we suggest that you create only one, using the page size that the majority of your tables use.
- User-temporary table spaces are used to store declared global temporary tables that store application-temporary data. User-temporary table spaces are not created by default at database creation time.

Advantages of an SMS table space

Advantages of an SMS table space are as follows:

- ▶ Space is not allocated by the system until it is required.
- Creating a database requires less initial work, because you do not have to predefine containers.
- ► A container is a physical storage device and is assigned to a table space.
- A single table space can span many containers, but each container can belong to only one table space.

Advantages of a DMS table space

Advantages of a DMS table space are as follows:

- ► The size of a table space can be increased by adding containers. Existing data is automatically rebalanced across the new set of containers to retain optimal I/O efficiency.
- ► A table can be split across multiple table spaces, based on the type of data being stored:
 - Long field data
 - Indexes
 - Regular table data

You might want to separate your table data for performance reasons, or to increase the amount of data stored for a table.

For example, you can have a table with 64 GB of regular table data, 64 GB of index data, and 2 TB of long data.

If you use 8 KB pages, the table data and the index data can be as much as 128 GB.

If you use 16 KB pages, the table data and the index data can be as much as 256 GB.

If you use 32 KB pages, the table data and the index data can be as much as 512 GB.

- ► The location of the data on the disk can be controlled, if this is allowed by the operating system.
- ▶ If all table data is in a single table space, a table space can be dropped and redefined with less overhead than dropping and redefining a table.
- ▶ In general, a well-tuned set of DMS table spaces will outperform SMS table spaces.

In general, small personal databases are easiest to manage with SMS table spaces. However, for large, growing databases, you will probably want to use only SMS table spaces for the temporary table spaces, and separate DMS table spaces, with multiple containers, for each table. In addition, you will probably want to store long field data and indexes in their own table spaces.



В

Worksheets

In this appendix, we provide worksheets that are intended for you to use during planning and installation of the Tivoli Storage Productivity Center. The worksheets are meant to be examples. Therefore you can decide whether you need to use them, for example, if you already have all or most of the information collected somewhere.

Tip: If the worksheets are too small for your handwriting, or you want to store the information in an electronic format, simply use a word processor or spreadsheet application, and use our examples as a guide, to create your own installation worksheets.

This appendix contains the following worksheets:

- ► User IDs and passwords:
 - Server Information
 - User IDs and passwords for key files and installation
 - LDAP Information
- Storage device information:
 - IBM System Storage Enterprise Storage Server, DS6000, DS8000
 - IBM System Storage DS3000, DS4000, DS5000 family products
 - IBM System Storage SAN Volume Controller (SVC)
 - IBM XIV Storage System (XIV)
 - IBM Storwize V7000 Midrange Disk System (Storwize V7000)
 - Other storage device information

User IDs and passwords

We created a table where you can record the user IDs and passwords that you need during the installation of Tivoli Storage Productivity Center, for reference during the installation of the components, and for future add-ons and agent deployment. Use this table for planning purposes.

You need one of the worksheets in the following sections for each computer where at least one component or agent of Tivoli Storage Productivity Center will be installed. The reason is because you can have multiple DB2 databases or logon accounts and you need to have the IDs of each DB2 individually.

Server information

Table B-1 contains information about the servers in the Tivoli Storage Productivity Center environment.

Table B-1 Tivoli Storage Productivity Center server

Server	Configuration information
Machine	
Host name	
IP address	··

In Table B-2, indicate (Yes or No) whether a manager or a component will be installed on this computer.

Table B-2 Managers or components installed

Manager or component	Installed (Y or N)?
Productivity Center for Disk	
Productivity Center for Replication	
Productivity Center for Data	
Tivoli Agent Manager	
DB2	

User IDs and passwords for key files and installation

Use Table B-3 to note the password that you used to lock the key file.

Table B-3 Password used to lock the key files

Default key file name	Key file name	Password
agentTrust.jks		
3		

Use Table B-4 to enter the user IDs and passwords that you used during the installation. Depending on the selected managers and components, certain lines are not used for this computer.

Table B-4 User IDs used on this machine

Element	Default or recommended user ID	Enter user ID	Enter password
DB2 DAS User	db2admin ^a		
DB2 Instance Owner	db2inst1		
DB2 Fenced User	db2fenc1		
Resource Manager	manager ^b		
Host Authentication			
Tivoli Storage Productivity Center Admin user ^c	tpcsuid ^a		
IBM WebSphere ^c			
Tivoli Storage Productivity Center for Replication Administrator ^c			

- a. This account can have any name you choose.
- b. This account name cannot be changed during the installation.
- c. If LDAP Authentication is selected, this value is overwritten.

LDAP information

If you plan to use an LDAP-compliant directory server for authentication, you are required to provide additional information during the installation of Tivoli Storage Productivity Center. Contact your LDAP administrator and gather the required information in Table B-5.

Table B-5 LDAP information

Element	Info	Actual value
LDAP Administrator contact info (name/email/phone)		
LDAP Server Type	e.g. Active Directory, OpenLDAP, IBM Tivoli Directory Server	
Load balancing configuration	See special precautions in LDAP	
LDAP Server Hostname		
LDAP Port Number	Default port = 389	
Secure LDAP Port Number	Default port = 636	
Secure LDAP Certificate file name		
Bind Distinguished Name (DN)		
Bind Password		
Relative DN for user names		
Attribute to use for user names	Default value = uid	
Relative DN for groups		
Attribute to use for groups	Default value = cn	
LDAP Tivoli Storage Productivity Center Administrator user name		

Element	Info	Actual value
LDAP Tivoli Storage Productivity Center Administrator password		
LDAP Tivoli Storage Productivity Center Administrator group		
backupConfig done?		

Storage device information

Use the worksheets in this section to gather important information about the storage devices that will be managed by Tivoli Storage Productivity Center. This information is necessary during configuration of the Tivoli Storage Productivity Center. You must have part of the information before you install the device-specific Common Object Model (CIM) agent, because this information sometimes depends on a specific code level.

Determine whether firewalls are in the IP path between the server or servers and the devices, which might not allow the necessary communication. In the first column of each table, enter as much information as possible to identify the devices later.

IBM System Storage Enterprise Storage Server, DS6000, DS8000

Use Table B-6 to collect the information about your IBM System Storage devices.

Important: Check the device support matrix for the associated CIM agent.

Table B-6 Enterprise Storage Server, DS6000, DS8000

Subsystem type, Name, location, organization	Both IP addresses	LIC ^a level	Username	Password	CIM agent host name and protocol

Subsystem type, Name, location, organization	Both IP addresses	LIC ^a level	Username	Password	CIM agent host name and protocol

a. Licensed Internal Code

IBM DS3000, DS4000, DS5000 family products

Use Table B-7 to collect the information about your DS3000, DS4000, DS5000 family products.

Table B-7 IBM DS3000, DS4000, DS5000 family products

Name, location, organization	Firmware level	IP address	CIM agent host name and protocol

IBM SAN Volume Controller (SVC)

Use Table B-8 to collect the information about your SVC devices.

Table B-8 SAN Volume Controller devices

Name, location, organization	Firmware level	IP address	Username	Password	CIM agent host name and protocol

IBM XIV Storage System (XIV)

Use Table B-8 to collect the information about your XIV devices.

Table B-9 XIV Storage System devices

Name, location, organization	Firmware level	IP address	Username	Password	CIM agent host name and protocol

IBM Storwize V7000 Midrange Disk System (Storwize)

Use Table B-8 to collect the information about your Storwize devices.

Table B-10 Storwize V7000 Midrange Disk System devices

Name, location, organization	Firmware level	IP address	Username	Password	CIM agent host name and protocol

Other storage device information

Use Table B-8 to collect the information about your other storage devices.

Table B-11 Other storage devices

Name, location, organization	Firmware level	IP address	Username	Password	CIM agent host name and protocol



C

Configuring X11 forwarding

In this appendix, we show the installation and configuration steps of tools that are used to achieve X11 forwarding in a firewalled environment. This information includes the installation of the prerequisite components on the AIX and Linux side, and the Windows workstation side. Following the instructions, you can use the graphical installers of the AIX and Linux distributions of Tivoli Storage Productivity Center V4.2 from your Windows workstation.

Preparing the display export

The various installers that are used to install the products described in this book use a graphical user interface (GUI) by default.

This section describes two methods; use whichever method suits you:

- ➤ X11 forwarding using a program such as Xming running on a Windows server (see "Preparation of the Windows workstation and Xming installation" on page 817)
- ▶ VNC server running on the AIX or Linux host system (see "VNC server" on page 831)

The **-silent** or **-console** options are not supported with Tivoli Storage Productivity Center V4.2.

The solution to achieve a display export described here is one of many possible ways to achieve the display export. Our servers and the environment we use are behind a firewall. It does not allow connections to be made from the AIX server behind the firewall to the outside computers in front of the firewall. Therefore we decided to implement the solution, that is based on the use of ssh, ssl, rpm, and Xming. The solution is described using an AIX server and a Windows workstation. It can also work with other UNIX distributions and Linux if the involved tools are applied properly.

Preparation of the AIX server

To install various tools for AIX, we use the **rpm** tool. Most Linux distributions already have **rpm** preinstalled; you have to separately install it for AIX.

Components are as follows:

► The rpm tool

The **rpm** tool for AIX is part of the AIX Toolbox for Linux Applications. It contains open source packages available for installation on IBM AIX 5L[™]. You can find the **rpm** tool and more information about the AIX Toolbox for Linux Applications at the following address:

http://www.ibm.com/systems/power/software/aix/linux/

You can also download the rpm tool directly from the following site:

ftp://public.dhe.ibm.com/aix/freeSoftware/aixtoolbox/INSTALLP/ppc/rpm.rte

After download, install the rpm tool by using smitty or installp.

► The Openss1 package

Also part of the AIX Toolbox for Linux Applications, but within the "Cryptographic Content for AIX" section, is the OpenSSL package for AIX (Figure C-1). To download cryptographic content, you have to log on to the download site with your IBM ID.

If you do not have an IBM ID, you can apply for one at the following site, at no cost:

https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=aixbp



Figure C-1 OpenSSL menu

▶ openssh

The third component that is used on AIX in this solution is the *OpenSSH on AIX* package. It is available on the open source software website:

https://sourceforge.net/projects/openssh-aix/

Access the latest version of OpenSSH on AIX and download the package.

After download, install the *OpenSSH* on *AIX* package using smitty / installp.

Preparation of the Windows workstation and Xming installation

Xming is an X Windows server for Windows. We use Xming to display the X Windows graphical interface, connecting to the AIX server by using ssh and enabling X11 forwarding. It is available for download at:

http://sourceforge.net/projects/xming/

Download the setup.exe file and install Xming as follows:

 On the Windows workstation you want to use to receive the X11 forwarding, double-click the installer file (Xming-6-9-0-31-setup.exe in our case). The Xming setup starts and welcomes you (Figure C-2). Click Next.



Figure C-2 Xming setup wizard

2. A new window opens (Figure C-3). It shows the directory where **Xming** will be installed. You can use the default or select another directory. Click **Next**.

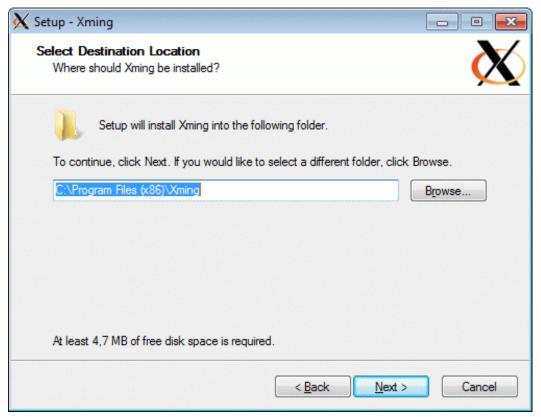


Figure C-3 Xming directory

 In the next window, select the software components to be installed. Keep the defaults as shown in Figure C-4. Make sure that Normal PuTTY Link SSH client or Portable PuTTY Link SSH client is selected. One of these will be necessary to enable X11 forwarding. Click Next to continue.

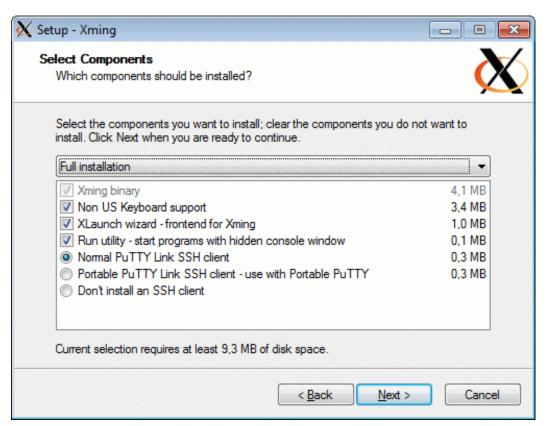


Figure C-4 Xming component selection

4. The next dialog offers to create a Start menu folder and shortcuts (Figure C-5). Click **Next** to continue.



Figure C-5 Xming Start Menu folder

5. In the next window (Figure C-6), select additional shortcuts to be created. You can keep the defaults. Click **Next** to continue.

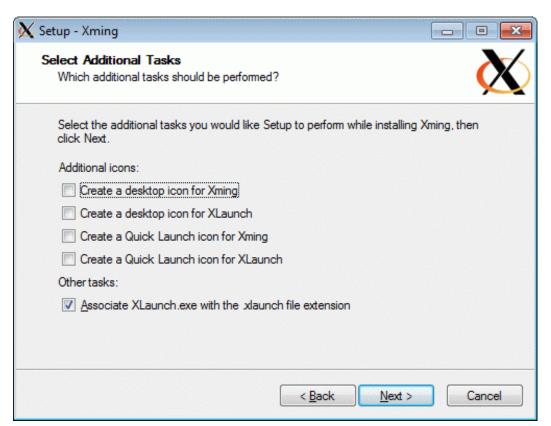


Figure C-6 Xming additional installer tasks

6. Before the installation begins, review your selections (Figure C-7) from the previous steps. Click **Install** to continue.

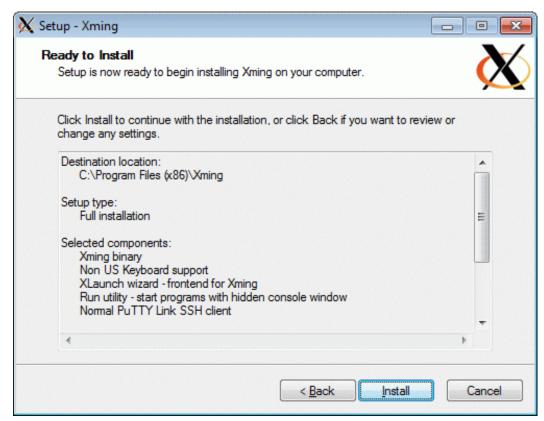


Figure C-7 Xming installation summary

7. A progress bar shows the installation status (Figure C-8).

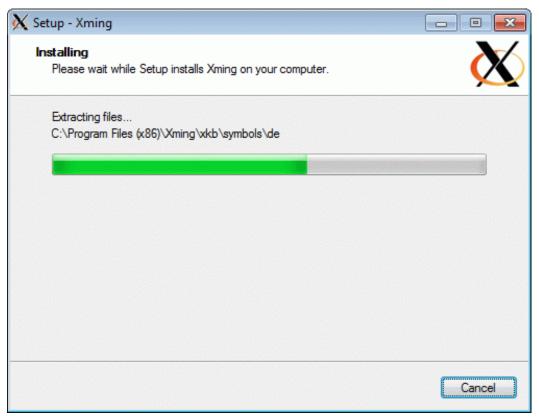


Figure C-8 Xming installation progress

8. When the installation is complete, a confirmation window opens (Figure C-9). Uncheck the "Launch Xming" check box or otherwise you need to stop the Xming later on for your newly created shortcut created in Figure C-14 on page 829 to work. Click **Finish** to exit he installer.



Figure C-9 Xming setup completed

Launching an Xming X Window session

Complete the following steps to launch a session:

 Log in to your AIX server, and review the /etc/ssh/sshd_config file to make sure that the following lines are present and are not commented out:

```
X11Forwarding yes
X11DisplayOffset 10
X11UseLocalhost yes
```

If the lines are not present, add them, and then restart the ssh daemon:

a. First you will need to determine the process number of the sshd service, issue:

```
ps -ef | grep sshd
```

b. Carefully browse through the results of the command and identify the ssh daemon.

```
[root@tpcblade6-11 ~] # ps -ef | grep sshd

root 4470 1 0 Sep09 ? 00:00:00 /usr/sbin/sshd

root 10599 4470 0 18:49 ? 00:00:00 sshd: root@pts/3

root 10953 10607 0 18:51 pts/3 00:00:00 grep sshd
```

c. In our example, the **ssh** daemon has the process number 4470. To restart **ssh** daemon, issue the following command:

kill -HUP 4470

You might have to disconnect your **ssh** session now and reconnect to enable the new settings for your session.

- 2. On the Windows workstation, start Xming by clicking Start → All Programs → Xming → Xlaunch to start the Xming configuration dialog.
- 3. A display settings panel opens (Figure C-10). Select **Multiple Windows** and click **Next** to continue.

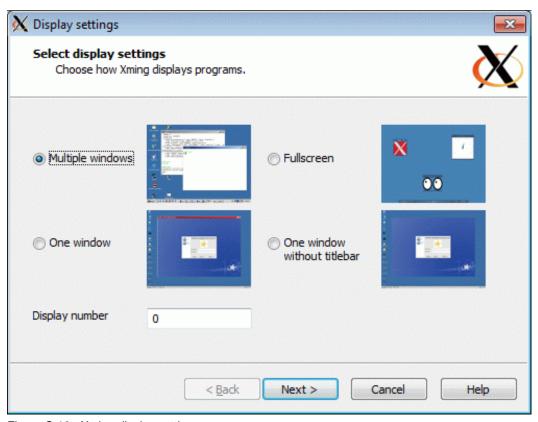


Figure C-10 Xming display settings

4. In the Session type window (Figure C-11), select Start a program. Click Next to continue.

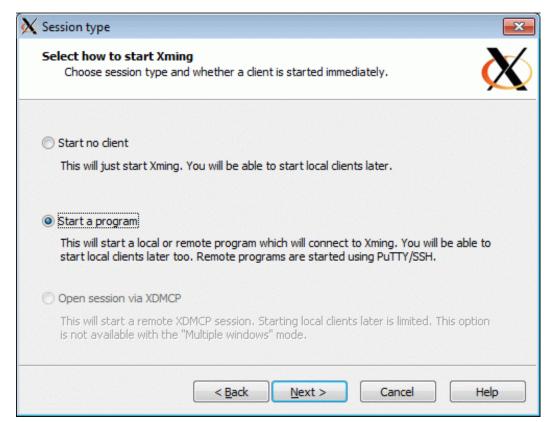


Figure C-11 Xming session type

5. In the next window (Figure C-12), select **xterm** as the program to be started, and for Run Remote, select **Using Putty (plink.exe)**. Enter your AIX host name, the user you will log in as (typically **root**) and the user password. Click **Next** to continue.

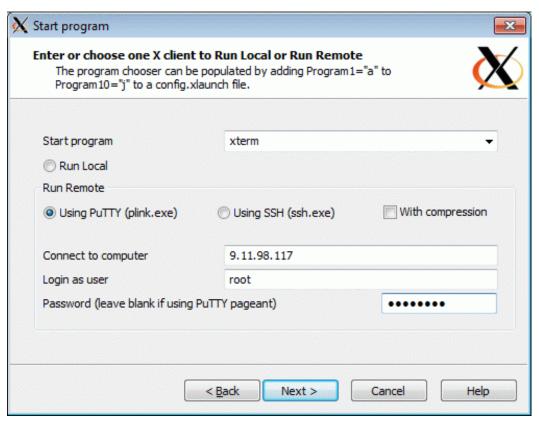


Figure C-12 Xming start program configuration

6. In the next window, you can specify additional settings and enable clipboard capability for your X Windows environment. You can keep the defaults (Figure C-13). Click **Next**.

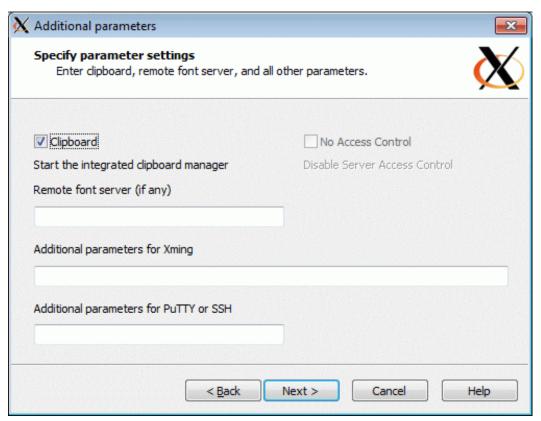


Figure C-13 Xming parameter settings

7. We suggest that you save the connection settings to use them in future sessions (Figure C-14). If you decide to do so, you can also save your password for an automatic login the next time you launch this connection. Click **Finish** to finalize your setup.

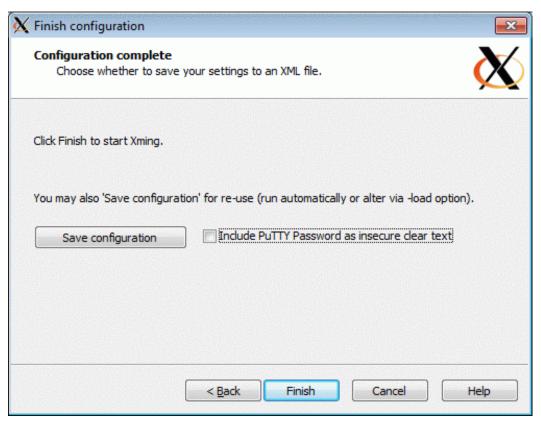


Figure C-14 Xming configuration save options

8. You have now successfully configured the Xming environment. To actually open up the terminal window from your AIX or Linux machine, you need to make sure that the Xming server is in the right state. To use your shortcut from Figure C-14 the Xming server needs to be stopped as it will be automatically invoked while running the shortcut. If it is started already, stop it by right-clicking XMing's tray icon shown at the top left of Figure C-15.



Figure C-15 Xming try icon

If you plan to manually invoke AIX or Linux programs that you want exported, start the Xming server manually by clicking $Start \rightarrow AII \ Programs \rightarrow Xming \rightarrow XMing$.

Open your saved connection from step 7 and an xterm window opens (Figure C-16).
 Congratulations, you can now launch graphical programs on your AIX server from your Windows workstation.

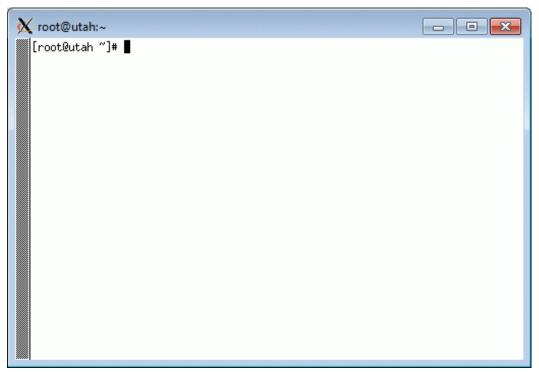


Figure C-16 xterm window exported via Xming

VNC server

If you prefer not to use programs such as Xming running on the Windows server, you can install and configure a VNC server on your AIX or Linux machine. VNC is a remote control software solution that allows you to view and fully interact with one computer desktop (the "VNC server") using a simple program (the "VNC viewer") or a current web browser with Java enabled on another computer over LAN.

Tip: Depending on the quality of your IP connection, you might be able to achieve better results using Xming or VNC. Connections with high speed, bandwidth, and low latency should allow you to have the best results with Xming. Connections with low speed, bandwidth, and high latency are better suited for VNC. With VNC you will also be able to reconnect to your session with all your programs still running in case your connection to the VNC server happens to drop.

We describe the necessary steps.

The VNC server for AIX is available for download from the IBM AIX Toolbox for Linux site:

http://www.ibm.com/systems/power/software/aix/linux/toolbox/alpha.html

Use the following installation steps:

- 1. Download the VNC RPM package from the AIX Toolbox site (at the time of this writing, the file name is vnc-3.3.3r2-6.aix5.1.ppc.rpm)
- 1. Install the RPM package (on the AIX server) using the following command:

```
rpm -Uhv vnc-3.3.3r2-6.aix5.1.ppc.rpm
```

- 2. Define a password for VNC access by using the **vncserver** command.
- 3. Start the VNC service by using the **vncserver** command as follows:

```
tpcblade4-14v3> vncserver
```

```
New 'X' desktop is tpcblade4-14v3:1
```

Starting applications specified in /home/root/.vnc/xstartup Log file is /home/root/.vnc/tpcblade4-14v3:1.log

4. Open your web browser and enter the name or IP of your AIX server, and port number 580X, where "X" is your assigned display number (from step 3). In our case, the VNC display is :1, so we use port 5801.

5. Log on using the password created in step 4. See Figure C-17.

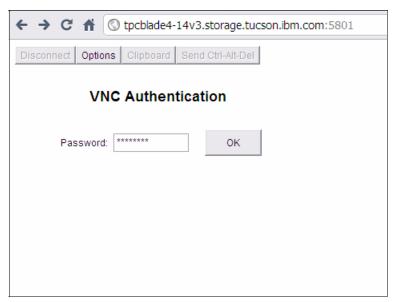


Figure C-17 VNC authentication

6. Obtain an X session with a terminal console, as shown in Figure C-18. Now you are ready to launch graphical applications on your AIX server.

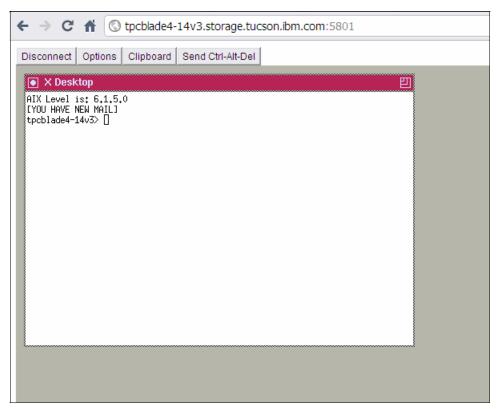


Figure C-18 VNC Viewer



D

LDAP considerations

This appendix provides additional information to help with understanding and implementing LDAP authentication in Tivoli Storage Productivity Center 4.2.

Overview

The Tivoli Storage Productivity Center 4.2 supports the Lightweight Directory Access Protocol (LDAP) as an alternative to OS authentication.

LDAP server implementations vary, and include but are not limited to these:

- Windows Active Directory
- ► IBM Tivoli Directory Server (ITDS)
- ▶ OpenLDAP

Only the last two of these are "true" implementations of the full LDAP protocol. Windows Active Directory was designed to support administration of computers and users in a Windows domain. A Windows Active Directory domain controller can serve as a LDAP server, but does not have the flexibility and configurability of true LDAP servers such as the IBM Tivoli Directory Server.

Organization

In addition, the organization of user entries in your LDAP directory can vary. It is important to understand the structure when considering how best to implement and assign user accounts for use with Tivoli Storage Productivity Center.

LDAP directory entries are implemented in a hierarchy, and the complexity and granularity vary from one installation to the next. The most important consideration is the entry you specify that will be used as the base for a downward search of the directory tree for all user accounts and groups to be used with Tivoli Storage Productivity Center. If this entry is not set at a point high enough in the tree that is common to all of the user accounts and groups to be used with Tivoli Storage Productivity Center, a directory search might not find a user account or group and authentication will fail.

Tivoli Integrated Portal and Tivoli Storage Productivity Center

The Tivoli Integrated Portal (TIP) is used to launch Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication. It supports a single sign on, which authenticates the user one time through the TIP login. This establishes credentials which are passed to the other Tivoli tools as they are launched, so you do not have to provide a user account and password for each tool.

Tip: Tivoli Common Reporting uses its own TIP instance, and so is not part of the single sign on for Tivoli Storage Productivity Center 4.2.

The authentication mechanism works through and is synchronized with the three WebSphere server instances that are part of Tivoli Storage Productivity Center. Changes to the authentication configuration must be propagated from the Tivoli Integrated Portal, to the Tivoli Storage Productivity Center Device Server, and to the Tivoli Storage Productivity Center for Replication Server.

Preparing to implement LDAP

A worksheet for LDAP information is provided in "LDAP information" on page 806. You should complete this worksheet before configuring LDAP.

LDAP can be configured for Tivoli Storage Productivity Center in three scenarios:

- 1. At initial installation of a new Tivoli Storage Productivity Center
- 2. Converting from OS authentication to LDAP
- 3. Converting from non-secure standard LDAP to secure LDAP

The Tivoli Storage Productivity Center InfoCenter has detailed instructions for each of these scenarios. Also see the Tivoli Storage Productivity Center support website:

http://www-947.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition

Also see 20.5, "Tivoli Storage Productivity Center support matrix" on page 792 for an example of how an end-user can search for a topic such as LDAP.

Your LDAP administrator will need to provide much of the information you will need to configure Tivoli Storage Productivity Center for LDAP. They should help you complete the LDAP worksheet, and provide answers to questions such as:

- What type of LDAP server do you have Windows AD? IBM Tivoli Directory Server(ITDS)? OpenLDAP?
- ▶ Do you have a single server, or do you use multiple servers for load balancing?
- ▶ Do you have a bind user and password or do you allow anonymous bind?
- ► If I want to configure secure LDAP, is there a special certificate file that should be used, or can I just retrieve the certificate from server port 636?

Important: before making any changes to the authentication configuration, you should back up the existing authentication configuration. The steps are documented in the Tivoli Storage Productivity Center InfoCenter, but we cover them briefly here because it is a vital step to provide a recovery point if a problem happens during reconfiguration.

Deciding whether to install with OS or with LDAP

As mentioned above, you can choose to configure LDAP when you install Tivoli Storage Productivity Center, or you can install with OS authentication and convert to LDAP later. What are the tradeoffs?

Installing Tivoli Storage Productivity Center with LDAP is easier than converting from OS after install. However:

- ► You cannot configure secure LDAP at install time, only non-secure standard LDAP. If the ultimate goal is to use secure LDAP, you will still have to go through much of the same conversion procedure and TIP panels as you would converting from OS to LDAP.
- ► If there are problems with your LDAP environment, you do not have an easy recovery option to fall back to OS authentication.

For this reason, it is best to install Tivoli Storage Productivity Center with OS authentication initially. The next section shows how you can back up this configuration for recovery in preparation for implementing LDAP.

backupConfig and restoreConfig

The 'backupConfig' and 'restoreConfig' commands are part of WebSphere. On Windows, the commands have a '.bat' extension, and on AIX and Linux they have a '.sh' extension. The 'backupConfig' command creates a .zip file backup of your current working authentication configuration before making changes, and the 'restoreConfig' command does what you would expect - it allows you to restore a working configuration from a backup.

In addition to 'backupConfig' and 'restoreConfig', there is a properties file that needs to be copied to a backup file. This file is named 'soap.client.props'.

We suggest that you copy the file, adding the string '.backup_yyyy-mm-dd' to the name of the copied file, to match the naming convention of the .zip file created by 'backupConfig':

Windows: copy soap.client.props soap.client.props.backup_2010-01-18

AIX/Linux: cp soap.client.props soap.client.props.backup_2010-01-18

Because Tivoli Storage Productivity Center has three embedded WebSphere instances, these steps must be performed for each instance. Here are the locations of the command and properties files ($x = C:\Pr$ or Windows, x = opt on AIX or Linux):

1. backupConfig (also restoreConfig) command location:

x/IBM/Tivoli/tip/profiles/TIPProfile/bin/backupConfig(.bat or .sh)

x/IBM/TPC/device/apps/was/profiles/deviceServer/bin/backupConfig(.bat or .sh)

x/IBM/replication/eWAS/profiles/CSM/bin/backupConfig(.bat or .sh)

2. soap.client.props file location:

x/IBM/Tivoli/tip/profiles/TIPProfile/properties/soap.client.props
x/IBM/TPC/device/apps/was/profiles/deviceServer/properties/soap.client.props
x/IBM/replication/eWAS/profiles/CSM/properties/soap.client.props

The 'backupConfig' command creates a .zip file named with the date/time stamp when the backup is taken. Figure D-1 displays a Windows example showing how this command is run for the TIP WebSphere instance. The procedure is similar for the device server and for the replication server. Note the name of the backup file that is created:

```
C:\Program Files\IBM\tivoli\tip\profiles\TIPProfile\bin>backupConfig -help
Usage: backupConfig [backup_file] [-nostop] [-quiet] [-logfile <filename>]
           [-replacelog] [-trace] [-username <username>] [-password <password>]
           [-profileName <profile>] [-help]
C:\Program Files\IBM\tivoli\tip\profiles\TIPProfile\bin>backupConfig
ADMU0116I: Tool information is being logged in file C:\Program
           Files\IBM\Tivoli\tip\profiles\TIPProfile\logs\backupConfig.log
ADMU0128I: Starting tool with the TIPProfile profile
ADMU5001I: Backing up config directory C:\Program
           Files\IBM\Tivoli\tip\profiles\TIPProfile\config to file C:\Program
           Files\IBM\tivoli\tip\profiles\TIPProfile\bin\WebSphereConfig 2010-01-18.zip
ADMU0505I: Servers found in configuration:
ADMU0506I: Server name: server1
ADMU2010I: Stopping all server processes for node TIPNode
ADMU7702I: Because server1 is registered to run as a Windows Service, the
           request to stop this server will be completed by stopping the
           associated Windows Service.
ADMU5002I: 384 files successfully backed up
ADMU6009I: Processing complete.
C:\Program Files\IBM\tivoli\tip\profiles\TIPProfile\bin>dir *.zip
Volume in drive C has no label.
 Volume Serial Number is E80E-CC9D
 Directory of C:\Program Files\IBM\tivoli\tip\profiles\TIPProfile\bin
12/09/2009 11:33 AM
                             7,857,792 WebSphereConfig 2009-12-09.zip
01/18/2010 12:27 PM 7,857,792 WebSphereConfig_2010-01-18.zip
2 File(s) 15,715,584 bytes
              0 Dir(s) 2,738,585,600 bytes free
C:\Program Files\IBM\tivoli\tip\profiles\TIPProfile\bin
```

Figure D-1 backupConfig command, Windows example for TIP WebSphere instance

The 'backupConfig' command will stop the WebSphere server when it runs, which will require you to restart the service after the command has completed.

Stopping and restarting services: Order matters

When making changes to the authentication configuration for Tivoli Storage Productivity Center, it is very important to stop and start the three WebSphere server services for the components in a specific order. This is necessary to allow changes to propagate from the TIP instance and synchronize with both the device server and the replication server instances.

Follow this order when stopping:

- 1. replication server (IBM WebSphere Application Server V6.1 CSM)
- 2. device server (IBM WebSphere Application Server V6.1 DeviceServer)
- 3. TIP (Tivoli Integrated Portal V2.2_TIPProfile_Port_16310)

When starting the services after a change, start them in this order:

- 1. TIP (Tivoli Integrated Portal V2.2_TIPProfile_Port_16310)
- 2. device server (IBM WebSphere Application Server V6.1 DeviceServer)
- 3. replication server (IBM WebSphere Application Server V6.1 CSM)

Attention: Manually stopping and starting services is preferred over rebooting the server. Allowing the services to start on their own after a reboot might introduce timing problems with synchronizing configuration changes, resulting in login failures. On Windows, you can set the services to 'manual' startup, reboot, then start the services in the order above and set them back to 'automatic' after the server comes up.

Troubleshooting authentication problems: Configuration files

In addition to the .zip file created by 'backupConfig' and the soap.client.props file, each WebSphere instance has two .xml configuration files that you can examine or send to support for help with resolving Tivoli Storage Productivity Center login or authentication problems.

- security.xml
- wimconfig.xml

These files are found in the following locations:

```
TIP (x/IBM/Tivoli/tip/...):

.../profiles/TIPProfile/config/cells/TIPCell/wim/config/wimconfig.xml

.../profiles/TIPProfile/config/cells/TIPCell/config/security.xml

.../profiles/TIPProfile/properties/soap.client.props

Device Server (x/IBM/TPC/device/...):

.../apps/was/profiles/deviceServer/config/cells/DefaultNode/wim/config/wimconfig.xml

.../apps/was/profiles/deviceServer/config/cells/DefaultNode/security.xml

.../apps/was/profiles/deviceServer/properties/soap.client.props

Replication Server (x/IBM/replication/...):

.../eWAS/profiles/CSM/config/cells/DefaultNode/wim/wimconfig.xml

.../eWAS/profiles/CSM/config/cells/security.xml

.../eWAS/profiles/CSM/config/cells/security.xml
```

Secure LDAP

The Tivoli Storage Productivity Center InfoCenter section on changing from OS authentication to LDAP also includes a link to instructions on configuring secure LDAP.

As a best practice, first configure standard (non-secure) LDAP, verify that it works correctly, and then back up the configuration before switching to secure LDAP.

As with the LDAP configuration procedure, discuss secure LDAP configuration with your LDAP admin. Make sure that secure LDAP is an available option, whether or not it can be used with anonymous bind, and also ask if there is a specific certificate file that you are required to use.

Useful LDAP tips

- ALWAYS back up your current working configuration -before- making any changes.
- ► It is a good idea to record the name of the .zip file created by 'backupConfig' and note the type of configuration it represents. Hint: you can rename the file to add an identifier to the front of the file name (such as 'os_' or 'ldap_') to clearly identify what configuration the backup represents.



Ε

Tivoli Common Reporting: Remote Installation

In this appendix, we cover the additional steps needed to install Tivoli Common Reporting on a remote server separate from the Tivoli Storage Productivity Center server.

Overview

Tivoli Common Reporting can be installed on a server other than the Tivoli Storage Productivity Center server. This option is preferable:

- ▶ When you have a large environment with a busy Tivoli Storage Productivity Center server
- ► When the Tivoli Storage Productivity Center server does not have enough physical resources (memory, disk space) to support the additional requirements for Tivoli Common Reporting

In order to install Tivoli Common Reporting on a remote server, you must install DB2 client software which provides the remote application the ability to connect to the Tivoli Storage Productivity Center server TPCDB database. The DB2 client is an option you can choose to install from the DB2 V9.7 package that is provided with Tivoli Storage Productivity Center.

After the DB2 client is installed, there are two DB2 commands that you run to configure the connection to the Tivoli Storage Productivity Center TPCDB database, then the steps to install and configure Tivoli Common Reporting are the same as given in section "TCR installation and setup" on page 606.

In the example we present here, we install the DB2 client and TCR on a Linux Red Hat 5.7 server with a connection to our Tivoli Storage Productivity Center V4.2.2 server TPCDB database running on Windows 2008.

Installing the DB2 client

The DB2 client is another option you can choose to install from the DB2 V9.7 installation package provided with Tivoli Storage Productivity Center. When you launch the DB2 installer, you can scroll to the bottom and choose an option to install the client as shown in Figure E-1.

Attention: The DB2 installer on Windows is slightly different, and presents only one option to install a DB2 client.

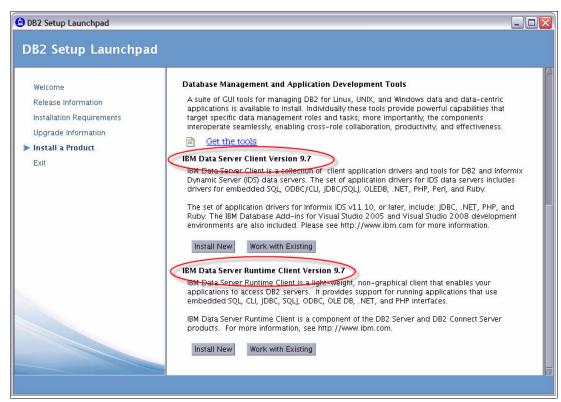


Figure E-1 DB2 installer on Linux, showing the two options that can be used to install the DB2 client

- ► The DB2 client installation is very straightforward.
- It is best to use a 'Typical installation, and you can safely take all of the defaults.

Tip: In our environment, there was already a DB2 install with a **db2inst1** instance. So we chose to use /opt/ibm/db2/V9.7_client as the install location to avoid collision with the existing DB2 V9.7 install, and chose to create a **db2inst2** instance for our client install as shown in Figure E-2 and Figure E-3.

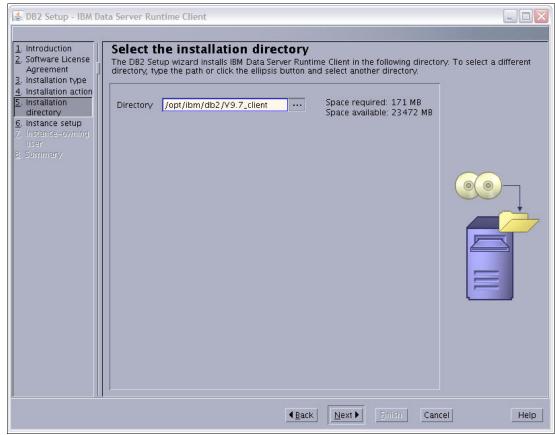


Figure E-2 Specifying a unique location for the DB2 client, necessary because another DB2 install was present

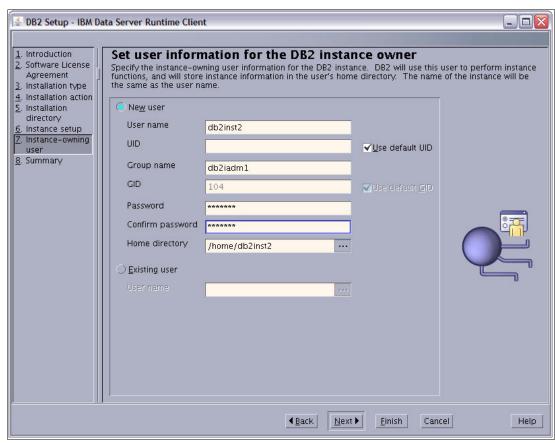


Figure E-3 Creating a new DB2 client instance 'db2inst2'

- 1. Complete the client installation.
- Do the remaining steps, which involve running two DB2 client commands to establish the connection to the TPCDB database on the Tivoli Storage Productivity Center server as shown in Example E-1.

Attention: Our environment uses a Windows TPC server called tpcblade3-13.

Example: E-1 Establish the connection to the TPCDB database

```
[root@madden07]# su - db2inst2
[db2inst2@madden07 ~]$ . sqllib/db2profile
[db2inst2@madden07 ~]$ db2
(c) Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.7.0

You can issue database manager commands and SQL statements from the command prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be the first few keywords of a database manager command. For example:
    ? CATALOG DATABASE for help on the CATALOG DATABASE command
    ? CATALOG for help on all of the CATALOG commands.
```

To exit db2 interactive mode, type QUIT at the command prompt. Outside interactive mode, all commands must be prefixed with 'db2'. To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => catalog tcpip node redbook remote tpcblade3-13.storage.tucson.ibm.com
ser

ver 50000 remote_instance db2 system tpcblade3-13 ostype win

DB20000I The CATALOG TCPIP NODE command completed successfully.

DB21056W Directory changes may not be effective until the directory cache is refreshed.

db2 => catalog database tpcdb as mytpcdb at node redbook authentication server DB20000I The CATALOG DATABASE command completed successfully.

DB21056W Directory changes may not be effective until the directory cache is refreshed.

db2 => terminate

DB20000I The TERMINATE command completed successfully.

[db2inst2@madden07 ~]\$

A brief explanation of the commands and parameters is in order:

► The first catalog command is long and wraps in our example. In Example E-2 we break it into its parts to make explaining it simpler.

Example: E-2 The first catalog command

CATALOG TCPIP NODE db2_node_name
REMOTE host_name
SERVER port_number
REMOTE_INSTANCE db2_instance_name
SYSTEM system_name
OSTYPE operating system type

where:

db2_node_name - is a new node you create in your DB2 client instance. In our example we chose 'redbook'.

host_name - this is the server name (or ip address) of the server where the TPCDB database is located. Normally this is your TPC server, and in our case this is 'tpcblade3-13.storage.tucson.ibm.com'.

port_number - the DB2 port that your TPC server is using, usually 50000.

db2_instance_name - the name of the DB2 instance on the TPC server. Since this
is a default DB2 install on Windows, the instance name is 'db2'.

system_name - the hostname of the TPC server where our TPCDB is located. We are using 'tpcblade3-13'.

operating_system_type - valid values for TCR are AIX, WIN, and LINUX. Our TPC server is Windows, therefore 'win' is what we use.

► Here is the second catalog command (Example E-3).

Example: E-3 The second catalog command

```
CATALOG DATABASE tpc_database_name AS tpc_database_alias_name
AT NODE db2_node_name AUTHENTICATION SERVER

where:
tpc_database_name - is your TPC database, normally TPCDB.

tpc_database_alias_name - this is a name you choose for your local client instance copy of the database. We chose 'mytpcdb'.

db2_node_name - the new node we created with the first command (ours was 'redbook').
```

With the DB2 client installed and configured, all that is needed is to install and configure TCR as described in "TCR installation and setup" on page 606.

References

For more information about DB2 client installation and configuration, see the following references in the DB2 InfoCenter.

```
http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp
http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.d
b2.luw.qb.client.doc/doc/t0005621.html
http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.d
b2.luw.qb.client.doc/doc/t0005622.html
```

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks publications

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ► IBM System Storage DS8000: LDAP Authentication, REDP-4505
- ► IBM Tivoli Storage Productivity Center V4.1 Release Guide, SG24-7725
- ► SAN Storage Performance Management Using Tivoli Storage Productivity Center, SG24-7364

You can search for, view, download, or order these documents and other Redbooks publications, Redpaper publications, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- ► IBM Tivoli Storage Productivity Center Storage Productivity Center for Replication, Version 4.2.2, Installation and Configuration Guide, SC27-2337-05
- ▶ IBM Tivoli Storage Productivity Center Storage Productivity Center for Replication, Version 4.2.2, User's Guide, SC27-2338-05
- ► IBM Tivoli Storage Productivity Center Storage Productivity Center for Replication, Version 4.2.2, Command-Line Interface Reference, SC27-2339-05
- ► IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference, SC27-2441

Online resources

These websites are also relevant as further information sources:

- ► IBM Tivoli Storage Productivity Center V4.2.2 Hints and Tips: https://www-304.ibm.com/support/docview.wss?uid=swg27008254
- ► Tivoli Storage Productivity Center support:

 $\verb|https://www.ibm.com/software/sysmgmt/products/support/IBMTotalStorageProductivityCenterStandardEdition.html|$

► IBM Software Support Lifecyle:

http://www.ibm.com/software/support/lifecycle/index_t.html

► Partner World Technical Delivery Assessment:

https://www-304.ibm.com/jct09002c/partnerworld/wps/servlet/ContentHandler/LLIE-6M7NYY/1c=en~US

► Tivoli Integrated Portal demonstration:

http://www14.software.ibm.com/webapp/download/demo.jsp?id=Tivoli+Integrated+Portal+Walkthrough+Aug08&locale=en

► XIV InfoCenter:

http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp

Help from IBM

IBM Support and downloads:

ibm.com/support

IBM Global Services:

ibm.com/services





IBM Tivoli Storage Productivity Center V4.2 Release Guide

(1.5" spine) 1.5"<-> 1.998" 789 <->1051 pages

Index

•	configure DB2 130
A	crfs command 134
Active Directory authentication 53	Custom installation 137
Active Directory Domain 273	DAS user 131
Active management server 372	Data Server Account Password 146
Add a CIFS share on NAS 301	Database administrator 145
Add Copy Sets 359	database repository location 131
wizard 412	database schema 135, 137, 145
Add Host List 271	database verification 141
Add NAS Server 307	DB2 communication 132
Add or configure a Brocade/McDATA Fabric 261	
Add or configure an IBM XIV 254	DB2 Enterprise Server Edition 129
Add Storage Resource Agents 269	DB2 installation 130 DB2 instance 131
administrative rights 32	
DB2 user 32	DB2 license 132
Administrator authority 265	DB2 SAMPLE database 132
Administrators group 265	DB2 user passwords 131
Advanced analytics 348	DB2 verification 132
Advanced Capacity Plan 452	DB2, table spaces 139
Advanced Tools 384	db2iadm1 group 131
AE daemon 352	db2profile script 136
startup 354	disable Tivoli Storage Productivity Center 169
agent	disk space 128
daemon service 216	ESE DB2 129
	gunzip command 135
installation 22	installation media 129, 134
issues on restore 685, 747, 781	LDAP authentication 146
performance improvement 341	NAS discovery 147
supported agents 6	prerequisites 128
upgrade scenarios 177, 212	progress indicator 151
Agent Assignment 340	RAM requirements 144
Agent code 297	restart DB2 132
Agent Manager 12, 176, 210	root account 142
backup 181	software prerequisites 128
agent migration	source instance profile 131
CLI 219	Storage Resource agent installation 144
errors 218	superuser role 146
success 219	Tivoli Integrated Portal instance 147
Agent upgrade through the CLI 212	Tivoli Integrated Portal ports 147
aggregates 694	Tivoli Storage Productivity Center for Replication 151
AIX 16, 266, 817	VNC viewer 135
chmod command 752	WebSphere Application Server ports 155
configuration manager 354	WebSphere user 146
defined sessions 356	X Window
GNU tar program 796	System display 135
Open HyperSwap 351	AIX VIO server 237
AIX 5.3	AIX VIOS 16
Storage Resource agent dependencies 266	alert
AIX 6.1	Alert panel, batch report 339
Storage Resource agent dependencies 266	notification 7
AIX 7.1 16	Allocated Volume Space 586
AIX configuration manager 353	analytics functions 2
AIX installation 126, 158	anonymous binds 54
Advanced security roles 146	APAR PK77578 56
authentication 148	API 542
command line 129	Arbitration Engine(AE) 353
	٠-٠ /

archive log files on file system 669, 732, 767 archive logging 711	capacity information 289, 291 Capacity licensing 588
AIX 711	Correlated definition 588
file system 759	Monitored computer space 588
parameter choices 658, 721, 757	Tivoli Storage Productivity Center for Data licensing
Assigned capacity 582	590
Asynchronous Global Mirror either direction with two-site practice 458	Tivoli Storage Productivity Center for Disk license 589
Asynchronous Global Mirror failover/failback 458	Tivoli Storage Productivity Center Standard Edition
Asynchronous Global Mirror single direction 458	licensing 592
authentication configuration 835	CD layout 23
authentication information 267	Certificate errors 786
auto-discovery function 233	Certificate Import Wizard 789
Automatic device discovery 348	cfgmgr command 383
	change management 348
autorun enabled 28	chmod command 752
Available File System Space 587	chunked capacity 581
Available Volume Space 585	CIM 227
average VDisk utilization 633	CIM agent 235, 261, 808 CIM client 241
	CIM indications 240
В	CIMOM 224, 559
back-end storage 346	Agent 558
background copy 402	CIMOM discovery 228, 234
backup	circular logging 711
cold 649	AIX 711
hot 649	Cisco 259, 784
Tivoli Storage Productivity Center server directories	Cisco Nexus 5000 Series 557
667	Cisco Nexus switches 556
backup considerations 667, 730, 766	CLI 568
backupConfig 836	Event Driven Jobs schedule 542
Base Port 53, 114, 147	cluster nodes 343
Basic Edition 9	Clustered scan support 343
batch report 333	Cognos 604
Alert panel 339	Content Database 613
definition of 568 type 334	cold backup 649, 710, 713, 749
**	collect Service Data 343
batch script AIX database backup 751	comma separated (CSV) file 387
create 751	comma separated values (CSV) 387, 568
BIRT 342, 570	file 568
Brocade	command-line interface 206, 212, 219
CIMOM configuration 257–258	CLI 568
Data Center Fabric Manager 13	Common agent 210
FCoE switch, 556	Common agent log 210
HBA 340	Common agent log 219 common agent registry 181
switches 556	Common Agent Services 209–210
Brocade 8000 556	Common Information Model 210
Brocade Data Center Fabric Manager 13, 558	Common Internet File System (CIFS) 266
Brocade DCFM 558	Common Reporting 619
Brocade DCFM 10.4 261	communication protocols 266
Brocade fabrics (FOS) 261	communication requests 266
Brocade Network Advisor 261	communication requirements 266
Brocade Server Management Console 559	component migration 178
Brocade Silk Worm 7800 (IBM 2005-R06 556	configuration changes 207
Browser back button 797	configuration History 568
Business Continuance Volume extended 581	configuration reports 568
С	Configure Devices Wizard data sources 237
candidate storage resource group 13, 438	Configure Devices wizard 11, 224, 236, 241-242

configure devices wizard		Configure Devices Wizard 237
CD wizard 237		data store 350
implementation 237		database 799
scenarios 236		attention required 34
Configure TCR Reports e-Mail 642		delete Tivoli Storage Manager archive logs 673
connection status values 413		repository 162
considerations for the Configure Devices wizard	237	repository growth 687
consistency group 376, 395, 410		views 569
naming convention 413		database backup
Consumed Volume Space 585-586		archive log files 669, 732, 767
Container Free Space 587		archive logging 658, 720, 756
Converged Enhanced Ethernet 556		archive logging to file system 661, 724, 760
Converged Network Adapter 557		automation 685
Copy Groups 406		delete Tivoli Storage Manager versions 672
copy service 404		delete versions 735, 770
copy set removal function 376		list of backup versions 675, 740, 775
Copy set soft removal 350, 376–377		managing versions 667, 730, 766
copy sets 17, 353, 360, 362, 378		offline backup 664, 710, 726, 762
Copy sets list 415		online backup 666, 710, 729, 764, 778
correlated definition 588		online backup script to file system 664, 726, 761
create a data source connection 617		online backup to file system 661, 666, 723, 730, 759
Create database schema option 40		765
Create local database option 43		
·		online Tivoli Storage Manager setup 658, 720, 756
Create NetApp Quota Job 330 Create session 356		parameter num_db_backups 650, 750
	004	rec_his_retentn parameter 650, 714, 750
Create Storage Resource Agent Upgrade panel	204	setup steps 650, 714, 750
Creating Diagnostic Packages 385		Tivoli Storage Manager password 656, 718, 754
credentials migration 192		to file system 652, 715, 751
credentials migration tool		database instance 5, 690
system credentials 171		error 47
crfs command 129		database register 283
CSV file 359, 387–388		database repository 206
CSV file restrictions 387		database restore 674, 739, 774
custom installation 23, 40, 80, 137, 184		agent issues 685, 747, 781
		circular logging 674, 739, 774
D		from file system 678–679, 742, 777
daemon communication 267		offline backup 675, 740, 775
daemon service 216, 266		online backup 669
daily history 696		roll forward 682, 745, 779
DAS user 131		Roll-Forward Pending 680, 744, 778
		Database Schema panel 187
Data agent 66		database scripts
legacy support 264		AIX 710
migration 176		database sizing
Data agents 12, 210, 237		Enterprise Storage Server 696
Data and Fabric agents 176		Fabric performance data 698
Data collection 247		Tivoli Storage Productivity Center for Data 699
data exposure 405		Database usage reports 291
data loss 418		database views 569
Data Management functions 348		Datapath View 296
Data Manager 297		DB2 352, 799
reporting		archive logging 658
Data Manager 599		backup image 675
data path explorer		backup on AIX
LUN correlation support 293		offline backup 751
Data Path View 296		online backup to file system 759
Data replication 415, 419		Tivoli Storage Manager API environment variables
Data server 4		753
port 9549 51, 112		backup on Linux
Data Server Account Password 52		offline backup 715
data sources 236		

online backup 720	buffer pool 571
backup on Windows	performance 43
Add new variables 653	sizing 43
dsmapipw command 656	DB2 database schema 568
Environment Variables 654	DB2 fix pack, 37
offline backup 657	DB2 installer 843
online backup to Tivoli Storage Manager 658	DB2 parameters 688
reboot 656	Failed log path 658, 721, 757
command line processor 659, 689	Primary log path 658, 721, 757
database backup on Linux 713	DB2 repository database
dsm.opt file 717	restore 674
dsm.sys file 718	DB2 UDB 82, 128
history file 714	DB2 versions 162
num_db_backups parameter 714	DB2ADMNS group, 42
Tivoli Storage Manager	db2adutl command 669, 732, 768
API environments variables 717	search backup image 676, 741, 776
database tuning 688	db2adutl query database command 673
datastore 350	db2sampl command 38, 99
db2_install 129	DCFM 261, 559
DMS table space	DCFM managed fabrics 565
advantages 801	decouple copy set 383
email notifications 96	Default Disk Jobs 541
fenced user 94	Default Fabric Jobs 541
full backup 658	default monitoring group 242
hints 796 installation 28, 128	deployment 341 Deployment Engine (DE 610
installation, verify 37, 46	Deployment Engine (DE) 610
instance owner 93	deployment job 265
license, AIX 132	deployment package 617
log files 46, 692	deployment schedules 265
offline backup script 657	Derby database 163
query 342	Detect storage events 348
restoring 674	device configuration wizard 222
setting variables 689	Device Server 5, 222, 838
setup wizard 86	port 9550. 51, 112
SMS table space	Device support matrix
advantages 800	associated CIM Agent 808
System managed space (SMS) 799	disable
System-temporary table spaces 800	Tivoli Storage Productivity Center 68, 166
table space 800	Tivoli Storage Productivity Center for Replication 164
tools catalog 33, 95	Disaster Recovery 351
tune the database manager 690	Disaster Recovery Planner 438, 457
user account 32	copy services session types 457
user passwords 131	Discover and probe the fabrics 439
User-temporary table spaces 800	Discover NAS Devices 314
variables setting 689	discover the NAS filer 307
DB2 9.7 12 DB2 Administration Server 32	Discovery Library Adapter (DLA) 8
DB2 backup	Discovery Library Adapter(DLA) 8 disk controllers 346
batch script 652	disk encryption 452
file, verify 674	Disk Management 348
DB2 client 842–843	Disk Manager 599
DB2 command-line processor window 650	Display PE Packages 386
DB2 commands 842	Distinct fabric discovery 340
db2adutl delete logs between 673	Distributed Computer Disk Space 591
db2adutl query database 673	DMS table space 800–801
db2adutl verify full taken 674	DNS suffix 25
list history backup all for 675	domain Administrator privileges 303
DB2 Control Center 46	domain controller information 273
DB2 database 5	domain name system 25

domain user 301	user-defined properties 784
double count storage 589	Fabric agent 12, 210
Download log packages 17	installation 66
down-tiering 606	migration 176
DR Planner 458	Fabric discovery job 562
DS3000 346	Fabric management 2
DS4000 346	fabric probe functions 340
DS6000 440	Fabrics and Switches
DS8000 403	user-defined properties 784
extent space efficient volumes 402	failover mechanism 226
failover 229	FC adapters 294
GUI 380	FCoE 556
HMC IP address 229	file system backup versions 667
NAPI configuration 251	file system Free Space chart
NAPI interaction 229	NAS 328
Open HyperSwap 353	File system reporting 326
redundant connection 229	FlashCopy 376, 457
SAN Planner 437	FlashCopy mapping 394, 401
single sign-on 5	FlashCopy operation 399, 401
storage 409	FlashCopy SE 402
volumes 403	FlashCopy session 394
Extent space efficient (ESE) 402	flat files offline backup 652
Track space efficient (TSE) 402	force command 379
DS8000 session enhancements 18, 350	-force option 279
DS8100 403	Formatted Capacity 588–589
DS8300 403	Formatted Space calculation 589
dsm.opt file 655, 717, 754	formatted storage space 580
dsm.sys file 718	free disk space 128
dsmapipw command 656, 718, 754–755	
DSopen API, NAPI interaction 235	G
	GEN0008E 176
E	GEN0324E 176
EIF probe 8	generate charts 574
embedded database 206	generate report 573
embedded repository 18, 206, 350	generate the Storwize V7000 internal disk report 599
Emulex CNA 556	Global Copy 376
Emulex HBAs 339	Global Mirror 19, 418
Enable NetBIOS over TCP/IP 27, 265	Global Mirror and Metro Mirror Path Manager 17
Encrypted Volumes 438	Global Mirror control paths 387
end-to-end planning 437	Global Mirror Data Exposure Graph 407
Enterprise Storage Server 440	Global Mirror Failover/Failback 15, 18, 410
Entities section 544	Global Will of Fallovel/Fallback 15, 10, 410
	Global Mirror session 17, 405
Entity Type 5/15	Global Mirror session 17, 405
Entity Type 545	multiple 409, 421
EPM 224	multiple 409, 421 planning 409, 421
EPM 224 probes 225	multiple 409, 421 planning 409, 421 practice 394
EPM 224 probes 225 ESE logical volume 403	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229 event changes 241	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239 GUI 4, 16, 422, 538, 568, 570
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229 event changes 241 Event Integration Facility 7	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239 GUI 4, 16, 422, 538, 568, 570 GUID 26
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229 event changes 241 Event Integration Facility 7 expansion units 346	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239 GUI 4, 16, 422, 538, 568, 570
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229 event changes 241 Event Integration Facility 7 expansion units 346 Extent space efficient (ESE) volumes 402	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239 GUI 4, 16, 422, 538, 568, 570 GUID 26 gunzip command 135
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229 event changes 241 Event Integration Facility 7 expansion units 346	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239 GUI 4, 16, 422, 538, 568, 570 GUID 26
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229 event changes 241 Event Integration Facility 7 expansion units 346 Extent space efficient (ESE) volumes 402 External Process Manager 224, 241	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239 GUI 4, 16, 422, 538, 568, 570 GUID 26 gunzip command 135
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229 event changes 241 Event Integration Facility 7 expansion units 346 Extent space efficient (ESE) volumes 402	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239 GUI 4, 16, 422, 538, 568, 570 GUID 26 gunzip command 135
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229 event changes 241 Event Integration Facility 7 expansion units 346 Extent space efficient (ESE) volumes 402 External Process Manager 224, 241	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239 GUI 4, 16, 422, 538, 568, 570 GUID 26 gunzip command 135
EPM 224 probes 225 ESE logical volume 403 ESE volumes 403 ESS/DS panel 388 ESSNI 11, 229 event changes 241 Event Integration Facility 7 expansion units 346 Extent space efficient (ESE) volumes 402 External Process Manager 224, 241	multiple 409, 421 planning 409, 421 practice 394 upgrade considerations 409 GNU tar program 796 graphical charts 568 group, definition of 239 GUI 4, 16, 422, 538, 568, 570 GUID 26 gunzip command 135 H hardware prerequisites 25, 127

reports 341-342	13
requirements 340	IBM TotalStorage Enterprise Storage Server (ESS) 294
health status, Storage Resource Group 441	IBM WebSphere 52
heterogeneous fabric 261	IBM WebSphere Application Server v6.1 - CSM 64
high availability configuration 372	IBM WebSphere Application Server v6.1 - Device Server
High level requirements for McDATA 258	64
High level requirements for QLogic 260	IBM XIV Storage System 18, 294
High-level requirements for Brocade 257	IBMCDB 180
High-level requirements for Cisco 259	Image Path 76
historical data 574	import quotas, NAS device 330
history aggregation 694	importing NetApp Quotas 330
HMC 229	in-band fabric interfaces 227
HMC connection 351	incremental FlashCopy 394–395
Hold I/O after Suspend 357	indication listener 240
Host Authentication Password 51	initial synchronization 207
Hosts and sessions 354	Install IBM Tivoli Common Reporting 608 Install Storage Resource agent 439
hot backup 649, 666, 710, 713, 729, 749, 764 HP 264	Install Storage Resource Agent on UNIX host 314
HP-UX 11i v3 16	installation
HP-UX native multipathing feature 439	authentication method 53
HTML 568	LDAP authentication 51
HyperSwap 350	licenses 40, 184
HyperSwap on System z 366	NAS discovery 52
hypervisor 333	planning disk space 128
Type Tibe 1	RAM warning 48
	remote database 50
I	sequence 38, 100
IBM ADE 168	timing 23
IBM ADE Service 170	Tivoli Storage Productivity Center for Replication 58
IBM Cognos Configuration Utility 642	verification 64
IBM Cognos Content Database 613	installation images 280
IBM DB2 Universal Database 282	installation wizard 195, 201, 212
IBM DS3000, DS4000, DS5000 series 294	InstallShield registries 180
IBM DS8000 440	integrated SMI Agent 559
IBM entry-level and midrange storage disk products 347 IBM High Availability Cluster Multiprocessing (IBM	internal disk reports
HACMP TM 343	Data Manager 599
IBM java runtime environment 606	Disk manager 599
IBM N series 297	Storwize V7000 598
configure windows Storage Resource agent 300	troubleshooting 603
IBM Power VM Virtual I/O Server 293	Internal Disks 599
IBM SDD driver 439	trouble shooting 603
IBM Software Lifecyle website 785	internal SSDs 598
IBM Storwize V7000 14	Internet Explorer address mismatch 787
SSH key 11	certificate conflict 786
IBM Storwize V7000 Midrange Disk System 294	interrupt upgrade 163
IBM System Director Storage Control 11	IO Group Memory 439
Version 6.2.1 11	IP address 804
IBM System Storage DS6000 series (DS6K) 294	IP network 317
IBM System Storage DS8000 series (DS8K 294	IP ports used from the native communication 233
IBM System Storage DS8000 series (DS8K) 294	IP ranges 229
IBM System Storage Productivity Center (SSPC) 10	IP Scan 228
IBM System Storage SAN Volume Controller (SVC) 5,	
294	1
IBM Tivoli Directory Server (ITDS) 834	J
IBM Tivoli Storage Productivity Center	Java 6 16
architecture overview 4	Java processes 182
IBM Tivoli Storage Productivity Center - Data Server 64	Job History 261, 540
IBM Tivoli Storage Productivity Center for Replication	Job History menu entry 542
17–18	job log 206
IBM Tivoli Storage Productivity Center Monitoring Agent	migration job 218

Job Management 262, 538	Tivoli Storage Productivity Center for Replication 118
job 539 run 539	verification 124 Web media 81, 127
schedule 538	WebSphere Application Server ports 120
Job Management panel 11, 233, 277–278, 457, 537	Linux Libraries
Context menu 542	commands 796
Entity Type 545	Tivoli Common Reporting 796
Icon bar 542	Linux libraries 796
Navigation Tree 542	Linux Red Hat 5.7 server 842
Runs and Jobs 546	local CLI installation 267
Schedules 545	Local Connection Status 412
Journal roles 392	Local Database Validation error 47
	local graphical installer 212, 267
K	Local or CLI installation of Storage Resource Agent 279
- -	Log File Viewer 542
Kill Process 77	log files
	CLI 65
L	Device server 65
launch-in-context 207	GUI 65
LDAP 806	Log packages 350, 384, 421
authentication 53, 112, 833	log space 661, 759
server, anonymous binds 54	logical paths 386–387
LDAP directory 834	logical subsystem (LSS 228
LDAP IBM RDN 55	lower performance storage tier 630
LDAP server 55	LPAR 295
LDAP server implementations 834	LSS 360, 387
LDAP worksheet 835	LUN Capacity 583
Legato NetWorker 711	LUN correlation support 294
license the NAS filer 305	
licensed capacity 588	M
Formatted capacity 588	Manage Exports 313
licensing methodology 346	Managed enclosures 346
licensing Standard Edition 592	Management Servers 373
licensing tab 314	managing backup versions 667, 730, 766
lightweight agent 264	Manually add NAS device 315
Lightweight Directory Access Protocol (LDAP) 834	manually entered NAS 304
Linux DM-Multipath driver 439	McDATA 258
Linux installation 79	McDATA EOS 558
component installation sequence 100	McDATA fabrics (EOS) 261
Data Server Account Password 113	Mdisk Group 621
database schema 101, 107	MDisk Group - VDisk Workload Activity Details Report
DB2 environment variables 99	621, 630
DB2 fenced user 94	MDisk groups 438, 605, 621, 624
DB2 instance owner 93	MDisk groups - Details Report 625
DB2 notification server 96	MDisk Groups - Workload Activity Report 621-622, 624
DB2 response file 89	MDisk groups - Workload activity Report 624
DB2 setup wizard 86	MDisks 438
DB2 tools catalog 95	media layout 23
db2iadm1 group access 98	memory requirement 186
db2prereqcheck 84	metadata 403
db2sampl command 99	metrics for capacity and capability 633
db2setup 84 df -h command 93	Metro Mirror 19, 376, 417
Host Authentication Password 113	Metro Mirror copy pairs 362
	Metro Mirror Failover/Failback 18–19, 410
mount point 101 NAS discovery 113	Metro Mirror Failover/Failback replication for DS8000
preregs 82	353
SA MP Base Component 90	Metro Mirror license 353
Tivoli Integrated Portal installation 117	Metro Mirror replication 17, 353, 417
Tivoli Integrated Portal Installation 117 Tivoli Integrated Portal ports 114	Metro Mirror session 351, 363
Tivon integrated Fortal ports 114	Open HyperSwap function 359

Metro Mirror session for data replication 418	N
Metro Mirror session type 350	N series 297
Metro Mirror sessions 418	Configure Unix Storage Resource Agent 313
Metro Mirror suspend policy 357	Models 299
Metro Mirror with Failover/Failback 15	monitoring 297
Microsoft	proxy agents 297
bginfo.exe 797	NAPI 220, 227
Microsoft Cluster Server 343	definition 228
Microsoft Excel 387	device discovery 228, 234
Microsoft SQL/Server 282	device ports 233
Midrange Edition 8	discovery 228
benefits 347	DS8000 229
supported devices 347–348	configuration 251
migrate the Data agent or Fabric agent to a Storage	DSopen API 235
Resource agent 213	EPM 224
Migrating from STAR tool 647	events 241
migration	migrate 220
multiple agents 215	migration
NAPI 220	CIMOM to NAPI 233
prerequisites 162	performance failover 226
Migration considerations 162	support
Migration from Tivoli Storage Productivity Center V3.x	2107 devices 227
178	IBM DS8000 227
Migration from Tivoli Storage Productivity Center V4.1	IBM SAN Volume Controller 227
178	IBM Storwize V7000 227
Migration methods 171	IBM XIV 227
migration tool	SVC recommendations 232
credentials 172	SVC versions 231
msvcr90.dll 172	tasks and panels 224
Mirror 410	when to use 228
modal dialog box 237	XIV
Monitored Computer Space 591	support 222
Monitored Computer Space report 590	NAPI device 220
Monitored Server Summary panel 311	NAPI discovery 224
monitored storage environment 568	NAS 297, 302
Monitoring Group 236, 247	Add Server 307
default 238	configure UNIX proxy 313
definition of 239	devices discovery 314
device assignment 237	discovery 52, 113
Subsystem Performance Monitor job 240	file system Free Space 328
Monitoring Group definition 238	import quotas 330
Monitoring Groups concept 240	manually add filer 306
monitoring jobs 568	map CIFS share 301
Monitoring Template 248	probe job 309
most active and least active Storage Resource Groups	Scan job 312
633	UNIX proxy Storage Resource agent 299
mount point 83, 129	Windows proxy agent 300
Mount the NAS root file-system to UNIX host 314	NAS device quotas 330
Mozilla Firefox	NAS devices 297, 317
certificate errors 786	Manage 302
MS Directory, add agents from 273	NAS devices information 298
msvcr90.dll 172	NAS filer 310, 322
m-type 558	asset information 325
Multi-Path Input/Output (MPIO) 352	license 315
multi-path planning 439	NAS FilerView 301, 313
Multiple Global Mirror sessions 18, 409, 421	NAS monitoring options 298
multiple NAS Servers 317	NAS Servers 308
multiple SAN views 348	NAS support
multiple volumes 450	SNMP queries 297
	NAS/NetApp Discovery job 304

Native API 221, 235	remove copy set 382
See NAPI	resume session 382
Native API (NAPI) 227	SDDPCM command 354, 365
Native API devices 220	session, prepared state 364
native interfaces 170	system setup 353
release level 171	unplanned 367
native storage system interfaces 11	Open HyperSwap replication 17, 350
DS8000 11	Open HyperSwap session 353, 365, 372
SVC 11	OpenLDAP 834
XIV 11	OpenSSH format 230
Navigating through the reports 641	OPENSSL package 816
Navigation Tree 539, 542	Operating system support 16
Job Management panel 539	Options panel 336
Navigation Tree Changes 222	Oracle table space 282
NetApp 300	organize backups on file system 668, 731, 766
NetApp Data ONTAP 323	orphan zone 265
NetApp devices 237	OS authentication 61, 835
NetApp filer mount 314	Out of Band Discovery 228
NetBIOS 26, 265	out-of-band fabric agents 11, 210
NetWare 264	
	overhead disk space 580
Network Advisor 558	
Configure a CIMOM Agent 561	P
Network Appliance devices 311	passphrases 230
Network Appliance filers 297	PASSWORDACCESS 655
Network File System (NFS) 297	Path Management wizard 389
New CLI commands for XIV Snapshot 416	Path Manager 350, 386–387
New performance metrics 14	port pairing 387
non-daemon agents 267	peer-to-peer remote copy (PPRC) 17
non-daemon service 266, 280	performance data 226
non-IBM storage 294	performance data 220
non-Select products 346	performance manager 14
Normal PuTTY Link SSH client 819	performance monitors 241, 439
	daily 696
0	· · · · · · · · · · · · · · · · · · ·
	history aggregation 695
offline backup 657, 675, 740, 775	Hourly 695
advantages 712	performance monitoring task 695
AIX 710, 751	performance report 568, 574–575
default method 711	performance reporting capabilities 347–348
Disadvantages 712	performance-based planning 439
disadvantages 712	performance-based virtual disk planning 439
flat files 652	Performing planned Open HyperSwap 368
scenario 652	Physical Address Extension (PAE) 797
Tivoli Storage Manager server 653, 752	physical media 24
offline database backups 664, 726, 762	plain text 568
OMNIbus 7	planned swap 366-367
Online backup	planner output 438
Advantages 712	platform dependencies 266
Disadvantage 712	plink.exe 827
online backup 712	Pool 410
AIX 710	port pairing 387
script, Windows 660	port pairs 388
online backup to file system 661, 666, 723, 730, 765	Portable PuTTY Link SSH client 819
Open HyperSwap 18, 350, 364, 366, 372	portpairings.csv file 388
decouple copy set 383	Power VM logical partition (LPAR) 293
DS8000 353	practice session types 410
function, Create session 356	practice sessions 394
MetroMirror session 363	predefined reports 568
overview 351	preinstallation tasks 47
planned swap 366-367	prerequisites for Open HyperSwap 351
properties panel 357	primary site 351

Private key 231	Replication product 5	
probe	replication server 838	
for multiple devices 225	report filters 572	
multiple device errors 250	report generation 571	
NAS filers 309	report type 571	
setup 285	reporting	
probe jobs 240	capacity 289, 291	
Process Explorer 74, 76	database views 569	
procexp64.exe file 75	Disk Manager 599	
Product code media layout and components 81	for databases 282	
Productivity Center Administrator role 265	GUI 568, 599	
Profiles tab 312	HBA 341	
Properties panel 357, 395	internal disk reports	
provision storage only 444	SVC 598	
provision storage with replication 444, 468	VDisk mirroring 594	
provision volumes 445	Reporting categories 569	
provisioning profile 459, 466	Reporting menu 568	
proxy agents 297, 308, 317	Reporting types 569	
Proxy DS CIMOMs 234	reports 568	
proxy SMI agents 261	assigned capacity 582	
Public key 231	available volume space 585	
PuTTY keys 230	back-end storage	
	volumes 581	
Q	consumed volume space 585–586	
Qlogic fabric 260	Database Capacity 588	
query Tivoli Storage Manager backups 670	DB2 query 342	
quota limit 330	disk space 581	
	double count storage 589	
n	formatted storage space 580	
R	licensed capacity Correlated definition 588	
RAID level 451		
RAID level selection 432	Network-wide Summary 292	
RAM	Storage Tier reports 604 Tier 0 579	
requirements 144	Tier 1 580	
warning message 48	Tier 2 581	
RDBMS Logins 283	Tier 3 582	
recovery site 417	Tier 5 585	
Red Hat Enterprise Linux Advanced Platform Version 5.4	Tier 6 586	
16 Red Hat Enterprise Linux Conver and Advanced Blotform	Tier 7 587	
Red Hat Enterprise Linux Server and Advanced Platform 6 16	unassigned capacity 582	
Red Hat Linux 5 266	unavailable disk space 581	
Redbooks website 849	Volume Space 584	
Contact us xix	reports based on Storage Resource Groups (SRG)	640
redundancy space 580	repository calculation templates 704	
register a database 283	request for price quotation (RPQ) 409	
registries backup 180	Requirements for NAPI 233	
Release I/O after Suspend 357	Resource Manager 210, 805	
remote database 50	response file 89	
remote deployment 267	restoreConfig 836	
Remote execution (REXEC) protocol 265	resume session 382	
Remote Installation 842	roll forward database restore 682, 745, 779	
remote job 343	Roll forward databases	
Remote shell protocol 265	to a point-in-time 684, 746, 780	
Remove Copy Sets wizard 378	to end of logs 683, 745, 779	
Replication Metro Mirror Failover/Failback session 353	Rollup Reports 176	
Replication Open HyperSwap 349	root account 142	
replication options 457	root authority 265	
replication planning 439, 457	root file-system 300	
replication planning with XIV 439	rpm tool 816	

RSH 265	FlashCopy 394
Rules for using CSV file 388	SAN Volume Controller incremental FlashCopy 17
Run a Probe job 309	SAN Volume Controller or Storwize V7000 consistency
	group 395
c	SAN Volume Controller session enhancements 17, 350
S	SAN Volume Controller Version 6.1 13
SA MP Base Component 90	Scan job 288, 321
SAN 439	NAS filer 312
SAN boot volume group 352	Scan/Probe Agent 308
SAN fabric information 210	Scheduling of TCR reports 644
SAN FS 264	schema log files 46
SAN plan 442	schema name 43
SAN Planner 12, 15, 437, 467	SDD 352
candidate storage resource group 13	SDD host error data 352
capacity plan 450	SDDPCM 352, 383
Disaster Recovery Planner 458	command 354, 365
DS8000 437	details 352
encrypted volumes 13	driver 353
generate recommendation 477	sddpcmgetdata 352
Metro Mirror with Failover/Failback 15	secondary site 351
new volumes 469, 497, 519	secure LDAP 839
output 479	Secure Shell (SSH) 230
provision storage 444	Secure Shell (SSH) protocol 265
purpose 438	Security roles 52
replication planning 440	security.xml 838
replication sessions 473	Select Provision Plan Details window 469
SAN Volume Controller (SVC) 437	SELECT statements 594
session properties 474	service collection script 343
Show Plan Topology View 478	Service Data SStorage Resource agent 343
Snapshot 15	Service Location Protocol (SLP) 235
space only mode 471	Session Details 413
Space Only planning 439	Simple Network Management Protocol (SNMP) 7
Space Only planning for all storage subsystems 12	single object 440
Space Only scenario 442	Single Sign-on 5
Space-efficient volumes	Single Sign-On feature 61
restrictions 12	Single-computer installation 611
storage provisioning 439	sizing performance collection for Enterprise Storage
Storwize V7000 437	Server 705
supported subsystems 440	sizing SAN Switch performance collection 707
SVC support 481, 494	Skip Workstations option 304
thin provisioning 452	SLES 11 339
VDisk create 486	SMI Agent 559
Vdisk mirroring 489	SMI-S 241
Wizard 443	SMS table space 800
XIV 437	Snapshot 15, 18, 410, 414, 457
zoning 465	Snapshot group 410
SAN Planner features 440	Snapshot session type 414
SAN Planner for IBM XIV	SNMP 227, 297–298
Asynchronous Global Mirror Failover/Failback 513	community name 304, 307, 316
Snapshot 494	MIB files 7
Synchronous Metro Mirror Failover/Failback 513	queries 297
SAN Planner recommendation 438	soap.client.props file 836
SAN Planner wizard 12, 438, 478	Solaris 9 266
SAN Volume Controller 346, 395	solid state disks 452
SSH key 11	source logical storage system 389
SAN Volume Controller (SVC) 392, 811	source logical subsystem (LSS) 386
SAN Volume Controller 6.1	source storage system 389
single sign-on 5	Space Only planning 438–439
SAN Volume Controller and Storwize V7000 front-end	Space Only workload profile 451
operations 438	Space Only workload profile option 437
SAN Volume Controller and Storwize V7000 incremental	

space-based planning. 12	benefits 210		
space-efficient volumes 392, 438	CLI install 279		
Specify Plan Content window 449	communication protocols 266		
SQL query, 38	daemon 66		
SQL statement 570, 596	daemon service 266		
SQL5005C error 796	Data agent migrate 211		
SRA 294	Data and Fabric agent functions 175		
SRGs 640	database monitoring 282		
SRM	deployment 268, 274, 276		
function 2	Fabric agent migrate 211		
legacy registry 181	fabric function 339		
SSH 230	force option 279		
key usage 231	graphical installer considerations 267		
session 253	IBM Power VM Virtual I/O Server 293		
sessions, SVC 230	install 47, 66, 110		
SSH daemon 824	installation parameters		
SSH key 230	installLoc 66		
SSH session 232	password 66		
SSL 266	serverIP 66		
SSPC	serverPort 66		
DB2 9.7 10	userID 66		
DS CIM agent management commands 10	introduction 263–264		
host bus adapter card 10	Job Management panel 278		
IBM Java 1.6 10	login local user 272		
IBM Storwize V7000 10	performance improvement 341		
IBM System Storage Productivity Center Version 1.5	register database 283		
10	resource utilization 210		
IBM System Storage SAN Volume Controller 6.1 10	Service Data 343		
IBM Tivoli Storage Productivity Center Basic Edition	Solaris 9 266		
4.2.1 10	supported platforms 266		
IBM Tivoli Storage Productivity Center for Replication	uninstall considerations 212		
Basic Edition 4.2.1 10	upgrade 200		
MC5 10	upgrade scenarios 211		
Microsoft Windows Server 2008 R2 10	VIOS 293		
System Storage DS8000 Release 6.0 10	VMware 331		
SSPC 1.5 10	Windows 272		
SSPC computer name 208	Windows Active Directory 273		
SSPC server 208	Storage Resource Agent fabric probe jobs 340		
Standard Edition 9	Storage Resource agent history 209		
licensing 592	Storage Resource agent installation 66		
standby management server 372	Storage Resource Agent not supported		
standby management server configuration 372	Filesystem Extension 264		
standby server 207	HP-UX on PA-RISC 264		
STAR report 647	NetWare 264		
STAR tool	Storage Resource Agent on VIOS 294		
Migrating from 647	Storage Resource Agent requirements 264		
Stopping Tivoli Integrated Portal 170	Storage Resource agents 11, 210, 331		
Storage Capacity Pyramid 578	Storage Resource Group 440, 467		
storage system capacity and utilization 578	alert overlay 440		
storage device 346	health status 441		
important information 807	VDisk Activity Report 632		
Storage device information 803, 807	VDisk Workload Ac 622		
storage devices or enclosures 346	VDisk Workload Activity Details Report 622		
storage provisioning 348, 439	Storage Resource Group - VDisk Activity Report 632		
Storage Resource agent 185, 202, 265–266, 297, 302,	Storage Resource Group - VDisk Workload Activity		
314, 340, 568	Details Report 622		
266	Storage Resource Group icon 441		
agent deploy 265	Storage Resource Group Management 458		
AIX 5.3 dependencies 266	Storage Resource Group panel 466		
AIX 6.1 266	Storage Resource Groups 440		

most active and least active 633 Storage Resource Groups (SRG) 437 storage resources 457 Storage Subsystem Credential Migration Tool panel storage subsystem credentials migration tool 170 Storage subsystem panel 224 storage system credential migration 175 storage system pairing 388 Storage Systems 412 Storage Tier reports 15, 567	SVC 252, 440 4.3 233 DB2 repository space 697 189 FlashCopy session 399 formatted space 580 incremental FlashCopy 394 replication CSV file 396 SAN Planner 437, 481 SAN PLanner support 481, 494 session username 253 space efficient volumes 392
back-end perspective 622 capability 605–606 capacity 605–606 down-tiering 605	space-efficient volumes 392–393 SSH key storage 231 user string 232 V4.2 233
e-Mail configuration 642 homogeneous storage pools 647 non-supported environments 647 Performance Monitor 639 performance monitor considerations 639	SVC and Storwize V7000 scenarios 252 SVC and Storwize V7000 space-efficient volumes 392 SVC virtual and mirrored disks 595 svcinfo Isuser 232 switches
Report types 621 scheduling 644 Storage Resource Groups (SRG) 640	toleration mode only support 13 Switches and Subsystem job 228 Sybase 282
storage virtualizer 605 TCR 604 TCR Reports List 621 Tiering 605	synchronized state 207 Synchronous Metro Mirror failover/failback 458 Synchronous Metro Mirror failover/failback with practice 458
TPCDB 604 up-tiering 605 Storage Tiering Activity Reporter (STAR) 15, 647 STORAGE_VOLUME_ID 597	Synchronous Metro Mirror single direction 458 system credentials 171 System managed (SMS) 43 System managed space (SMS 799
STORAGEPOOL 598 STORAGESUBSYSTEM view 596 Storwize V7000 xv, 231, 346, 392, 395, 440	System Storage DS8000 FlashCopy 395
C6.1 and later 233 incremental FlashCopy 394 internal disk report 599 managed-disk groups 5	T table space 800 tables of data 568
Native API 230 online storage optimization 14 SAN Planner 481	Takeover 372 takeover 372 Takeover action 374
single sign-on 5 special considerations 231 SSH keys 230 storage tier reports xv	Target Available state 367 target logical storage system 389 target LSS 386 target storage system 389 TBSM 8
VDisks 594 Version 5 and later 232 Storwize V7000 and SVC internal (local) disk reports 5 Storwize V7000 space-efficient volumes 392	TCR 606, 644, 842 architecture 605 command line tool 617
Structured Query Language (SQL 569 subnets 228 Subsystem Device Driver Path Control Module 352	create a data source connection 616 deployment 606 e-Mail configuration 642 Installation assistant 607–608
subsystem storage management 2 SUBSYSTEM_ID 596 Sun Solaris 9 and 10 16 Support Assistant 785	instance 615, 619 minimum system requirements 606 reports
support Assistant 765 support matrix 792 supported web browser 616 suspend action 380	scheduling 644 Reports List 621 TIP instance 612
suspend session 380	virtualizer boundary 641 TCR architecture 605

TCR deployment 606	Basic Edition 9			
TCR Installation assistant 608	Change DB2 active logs directory 692			
TCR instance 619	component install 38, 100			
Temporary table spaces 800	database schema 50			
thin provisioning 403	disable 166			
Three Site Business Continuity (BC) license 63	license 40, 184			
Three-site Metro Global Mirror 458	memory requirements for installation 186			
Three-site Metro Global Mirror with practice 458	Midrange Edition 8, 346			
thresholds 14	supported devices 346			
Tier 0	repository space 697			
Raw capacity of the storage subsystems 579	restart services 691			
Tier 1	server 807			
Formatted, raw lost to RAID 5, and unformatted	server stop AIX, AIX			
capacity 580	Tivoli Storage Productivity Center server stop 169			
Tier 2	services 64			
Chunked capacity and not allocated capacity 581	Standard edition 9			
Tier 3	start services 693			
Assigned and unassigned capacity 582	superuser 146			
Tier 4	upgrade			
LUN capacity as seen and not seen by operating	process explorer 74			
system 583	upgrade, database repository 162			
Tier 5	Tivoli Storage productivity Center			
Allocated and unallocated volume space 585	database sizing 693			
Tier 6	Tivoli Storage Productivity Center 4.2 Information Center			
Logical volume capacity and available logical volume	206			
capacity 586	Tivoli Storage Productivity Center agent			
Tier 7	upgrading 212			
calculation 587	Tivoli Storage Productivity Center dashboard 311			
File system total, and available file system and	Tivoli Storage Productivity Center Data agent 568			
database capacity 587	Tivoli Storage Productivity Center Data Manager for			
Tiering 605	Databases 282			
TIP 838	Tivoli Storage Productivity Center family 346			
TIP Administrator ID 53	Tivoli Storage Productivity Center for Data 8			
TIP login 834	database sizing 699			
Tivoli 126, 157, 841	repository tables 699			
Tivoli Application Dependency Discovery Manager	T_STAT_FILE table 701			
(TADDM) 8	Tivoli Storage Productivity Center for Data licensing 590			
Tivoli Business Service Manager (TBSM) 8	Monitored Computer Space report 590			
Tivoli Common Reporting 5, 606, 841–842	Tivoli Storage Productivity Center for Disk 8, 15			
browser back button 797	Tivoli Storage Productivity Center for Disk license 589			
Tivoli Common Reporting (TCR) 604	Tivoli Storage Productivity Center for Disk Midrange			
Tivoli Enterprise Console 7	Edition 346			
Tivoli Enterprise Console IBM Netcool OMNIbus 7	Tivoli Storage Productivity Center for Disk Select 8, 15,			
Tivoli GUID 667	346			
setting 181	Tivoli Storage productivity Center for Disk Select xv			
Tivoli Integrated Portal 5, 57, 158, 168, 208	Tivoli Storage Productivity Center for Replication 9, 61,			
one instance 147	81, 376, 404, 439			
ports 147	AE daemon 352			
stop 170	AIX modules 352			
Tivoli Integrated Portal (TIP) 834	copy set soft removal 376			
Tivoli Productivity Center for Replication 9	copy sets 362			
Tivoli Storage Manager 7	Create sessions 356			
API environments variables 653	datastore 350			
backup script, offline backup 657, 719, 755	decouple copy set 383			
DB2 parameters 658	define site locations 358			
option file 655, 717, 754	disable 68, 163			
password 656, 718, 754	force command 379			
server database backup 653, 752	high availability configuration 372			
Tivoli Storage Manager Backup-Archive client 716	install 58			
Tivoli Storage Productivity Center 803	Linux and AIX 165			

Log packages 384	Total File System Capacity and Free Space 587		
new function 350	Total Instance Storage Capacity 289		
Path Manager 386	total managed space 580		
port pairing 387	total raw disk space 580		
portpairings.csv file 388	TPC_backup_offline_file.bat script 665, 727, 762		
ports 62, 351	TPC_Reports.zip file 606		
repository upgrade 207	tpc_svc.pem 253		
resume session 382	TPCDB 180, 605, 842		
SDDPCM details 352	TPCDB database 43, 49, 107, 673		
server 166	logs 692		
service disabled 165	tune 690		
Standby server 372	TPCREPORT 7, 570		
standby server 207	TPCStorage sensor 8		
	TPCTOOL 7, 568, 570		
suspend session 380–381 takeover 372			
	track space efficient 405		
two and three Site Business Continuity license 9	logical volumes 402		
upgrade 163, 199	Track space efficient (TSE) volumes 402		
upgrade interrupt 163	Troubleshooting issues with Metro Mirror and Global		
Tivoli Storage Productivity Center for Replication Console	Mirror sessions 423		
421	trusted domain 302		
Tivoli Storage Productivity Center for Replication	TSE volumes 403		
database 362	TSRMsrv1 ID 52		
Tivoli Storage Productivity Center for Replication license	Two Site Business Continuity license 351		
439	Typical installation 22–23, 31, 40, 80, 184, 843		
Tivoli Storage Productivity Center for Replication server.	duration 81		
409			
Tivoli Storage Productivity Center for Replication	11		
web-user interface 63	U		
Tivoli Storage Productivity Center GUI 784	Unassigned capacity 582		
Tivoli Storage Productivity Center Hints and Tips 797	Unassigned Volume Space 583		
Tivoli Storage Productivity Center Navigation Tree 570	Unavailable Disk Space 580		
Tivoli Storage Productivity Center reporting 619	Unformatted Disk Space 580		
Tivoli Storage Productivity Center reports 15, 578	Unique sysName 299		
Tivoli Storage Productivity Center reports available with	UNIX Proxy Agent requirements 299		
TCR 620	UNIX proxy Storage Resource agent 299		
Tivoli Storage Productivity Center Select xv, 9, 14,	configure 313		
345–346	UNIX-attached NAS 306		
Tivoli Storage Productivity Center Standard Edition 15	unplanned swap 367		
Tivoli Storage Productivity Center Standard Edition	UPDATE DB CFG command 659, 663, 722, 725, 758,		
•	761		
licensing 592	Update Job Status 304		
Tivoli Storage Productivity Center Storage Tier reports	upgrade		
15	considerations 162		
Tivoli Storage Productivity Center support website	failure 180		
16	from V3.3.2 178		
Tivoli Storage Productivity Center Tier Reports 5, 604	Storage Resource agent 200		
Tivoli Storage Productivity Center TPCDB database 605	Upgrade Agents 204		
Tivoli Storage Productivity Center Two Site Business	upgrade interrupt 163		
Continuity 417	upgrade interrupt 103 upgrade job 205		
Tivoli Storage Productivity Center V4.1 568	. •		
Tivoli Storage Productivity Center V4.2	upgrade migration scenarios 177		
Job History 261	upgrade scenarios		
Tivoli Storage Productivity Center V4.2.2 Storage Tier	agents 177, 212		
reports 647	Storage Resource agent 211		
Tivoli Storage reports	Up-tiering 606, 635		
front-end perspective 627	UR (Uncommitted Read 596		
Topology View	user credentials 237		
Storage Resource Group 440	user ID 805		
Topology Viewer 295, 446, 568	user IDs and passwords 804		
NAS filer 323	user interfaces		
Topology Viewer selection for SAN Planner 446	UI 6		

V	Windows installation		
VCM database 581	DB2 reboot 47		
VDisk 457, 630	directory name 24		
VDisk based reports 636	media layout 23		
VDisk candidates for relocation 633	overview 22		
VDisk Details Report 622	prerequisites 24		
VDisk mirroring	Windows Internet Explorer		
Storwize V7000 481	Certificate errors 786		
SVC 481	certificate mismatch errors 787		
VDisk mirroring planning 439	messages 786		
VDisk mirroring reporting 594	notifications 786		
Storwize V7000 594	Windows Proxy Agent Requirements 300		
SVC 594	Windows registry 181		
VDisk Workload Activity Reports 622	Windows Resource Monitor 615		
VDISK_COPY view, 597	Windows server message block protocol (SMB) 2		
VDISK_COPY_TYPE 597	Windows Service Information 272		
VDisks 598	Windows-attached NAS 305		
VDisks that are mirrored 598	worksheets 803		
version installed 70	DSXXX 803		
View certificates 788	LDAP information 806		
View Log File 219	other Storage devices 814		
VIOS 293	server information 804		
VIOS LUN correlation 293	Storage Device information 807		
Virtual Disk Mirroring 489	Storwize devices. 813		
Virtual Disks 444	SVC devices 811		
virtual machine disk 333	User IDs and passwords 805		
Virtual SCSI 14	XIV devices 812		
VMware			
machine assets 331	X		
Storage Resource agent 331	X Window		
VMware guest images 339	server 135, 817		
VMware Hypervisor 237	X Window server application 135		
VNC 831	X Window System display 135		
VNC server 831	X Windows server 817		
VNC Viewer 135, 831	X11 forwarding 815		
Volume Configuration Management database 581	AIX 816		
Volume Group 365	Openssl 816		
volume ID 416	rpm tool 816		
volume nickname 416	Cygwin 816, 834, 842		
volume pairs 361	OPENSSL package 816		
Volume Space 584	Windows		
Volumes 444	Launch Xming" 824		
volumes panel 393 VSANs 784	Xming 817		
VSAINS 704	X11 graphical capability 82		
	XBSA 711		
W	XIV 222, 346, 429, 440		
WebSphere 836	Create Volume wizard 434		
WebSphere Application Server	Firmware levels 428		
admin ID and password 113	Global Mirror Failover/Failback 18		
WebSphere configuration files 180	interaction 233		
WebSphere Information 612	Metro Mirror Failover/Failback 18		
wimconfig.xml 838	NAPI support 222		
Windows	performance counters 428		
32-bit environments 797	performance statistics 429		
Windows 2008 265, 842	RAID level 432		
Windows 2008 R2 16	SAN Planner 437		
64-bit processor 10	Snapshot 18		
Windows 2008 Standard Edition 16	Storage Optimizer 435		
Windows firewall 265	storage provisioning 431		
	support 427		

three IP addresses 428 Tivoli Storage Productivity Center configuration 254 XIV connections 411 XIV consistency groups 413 XIV enhancements 410 XIV Gen2 hardware 410 XIV Metro Mirror session 417 XIV performance data 226 XIV session enhancements 350 Global Mirror Failover/Failback session support 350 Metro Mirror Failover/Failback session support 350 Snapshot session support 350 XIV storage system 13, 410, 413 XIV system configuration 419 XIV system enhancements 14 XIV system sessions 416 XIV system volume ID 416 Xming 816

Ζ

zoneset 265 zoning 438



IBM Tivoli Storage Productivity Center V4.2 Release Guide



Master the new capabilities of the Tivoli Storage Productivity Center

Integrate XIV and Storwize V7000 for storage resource management

Optimize your environment with Storage Tier reports

IBM Tivoli Storage Productivity Center V4.2 is a feature-rich storage management software suite. The integrated suite provides detailed monitoring, reporting, and management within a single console.

This IBM Redbooks publication is intended for storage administrators and users who are installing and using the features and functions in IBM Tivoli Storage Productivity Center V4.2. The information in the book can be used to plan for, install, and customize the components of Tivoli Storage Productivity Center in your storage infrastructure. This release contains several important functional enhancements:

- New licensing for Tivoli Storage Productivity Center Select and for Tivoli Storage productivity Center for Disk Select to provide full Tivoli Storage Productivity Center functionality to Midrange.
- Support for IBM Storwize V7000. Also new is the capability for customers to generate and see Storwize V7000 and SAN Volume Controller (SVC) Internal Disks report.
- ► New IBM Tivoli Tier reports that help clients make important decisions about storage tiering by utilizing estimated capability and actual performance data for IBM System Storage SAN Volume Controller and IBM Storwize V7000 storage tier reports.
- ► IBM XIV Storage System support has been updated, adding discovery, provisioning, and performance management, as well as new replication support for three new XIV sessions: Snapshot, Metro Mirror failover/failback, and Global Mirror Failover/Failback.

Step-by-step procedures are provided to help you with tasks such as migrating to Storage Resource agents, using Native APIs, using SAN configuration planning functions, and maintaining your IBM DB2 database repository.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information: ibm.com/redbooks

SG24-7894-01

ISBN 0738436275